**XPrize for Carbon Removal** › Finalists vie for $50 million from Elon Musk
P. 5

**Quantum Computing's Achilles' Heel** › Unavoidable errors and how to fix them
P. 28

**How the FCC Settles Radio Turf Wars** › When every party thinks it's right
P. 34

FOR THE TECHNOLOGY INSIDER
JULY 2022

# IEEE Spectrum

## Recycling Is a Mess. Can AI Sort It Out?

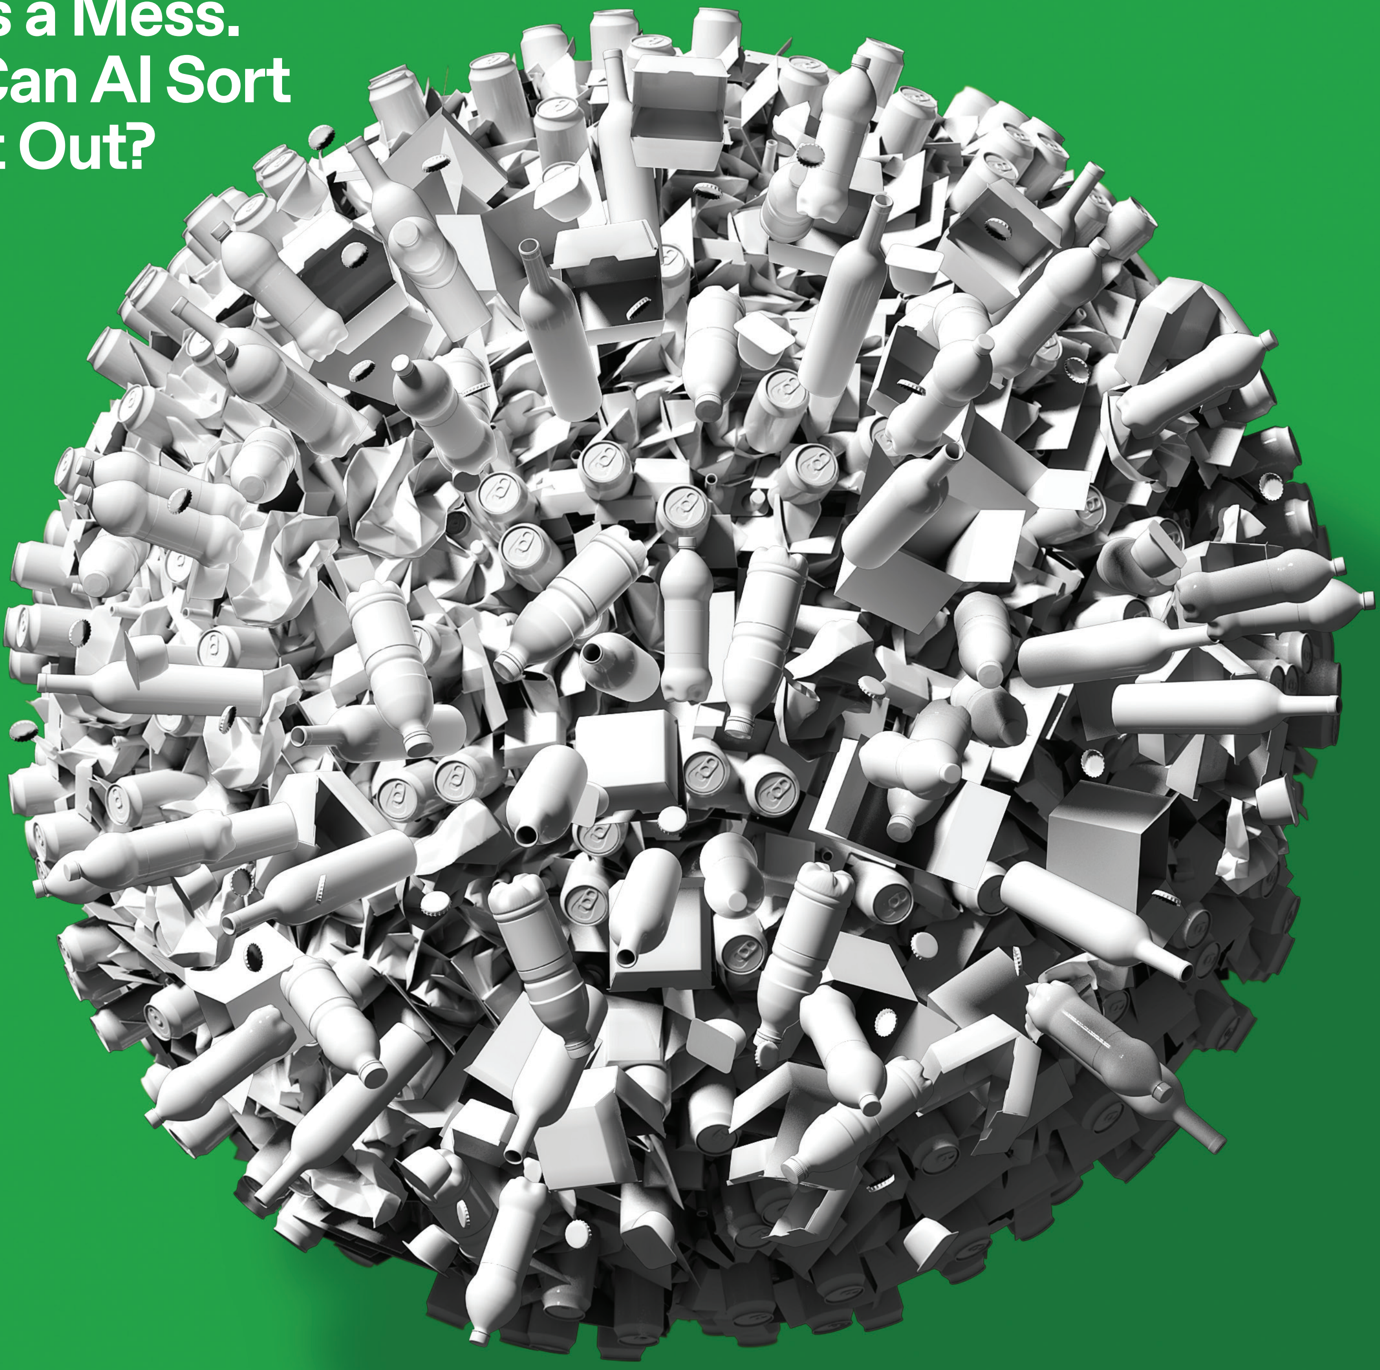

IEEE

VOLUME 59 / ISSUE 7 JULY 2022

# Contents

34

## Radio-Spectrum Turf Wars

Here's what typically ensues when interested parties can't agree.
**By Mitchell Lazarus**

LEFT: KEVIN KEY/SLWORKING/GETTY IMAGES; RIGHT: IIT MADRAS

## AI Takes a Dumpster Dive

22

Computer-vision systems learn to beat humans at the recycling game.
**By Jason Calaiaro**

## Quantum Error Correction at the Threshold

28

This great hurdle for quantum computing may soon be overcome.
**By Michael J. Biercuk & Thomas M. Stace**

## Weaning India From Coal

40

The move to 100 percent renewable energy starts with the commercial sector.
**By Ashok Jhunjhunwala, Kaushal Kumar Jha & Anson Sando**

ON THE COVER:
Illustration by Eddie Guy

BY HARRY GOLDSTEIN

# AI Can Help Make Recycling Better

## But only humans can solve the plastics problem

**Garbage is a** global problem that each of us contributes to. Since the 1970s, we've all been told we can help fix that problem by assiduously recycling bottles and cans, boxes and newspapers.

So far, though, we haven't been up to the task. Only 16 percent of the 2.1 billion tonnes of solid waste that the world produces every year gets recycled. The U.S. Environmental Protection Agency estimates that the United States recycled only about 32 percent of its garbage in 2018, putting the country in the middle of the pack worldwide. Germany, on the high end, captures about 65 percent, while Chile and Turkey barely do anything, recycling a mere 1 percent of their trash, according to a 2015 report by the Organization for Economic Cooperation and Development (OECD).

Here in the United States, of the 32 percent of the trash that we *try* to recycle, about 70 to 90 percent actually *gets* recycled, as Jason Calaiaro of Amp Robotics points out in “AI Takes a Dumpster Dive,” on page 22. The technology that Calaiaro's company is developing could move us closer to 100 percent. But it would have no effect on the two-thirds of the waste stream that never makes it to recyclers.

Certainly, the marginal gains realized by AI and robotics will help the bottom lines of recycling companies, making it profitable for them to recover more useful materials from waste. But to make a bigger difference, we need to address the problem at the beginning of the process: Manufacturers and packaging companies must shift to more sustainable designs that use less material or more recyclable ones.

According to the Joint Research Centre of the European Commission, more than “80 percent of all product-related environmental impacts are determined during the design phase of a product.”

According to the Joint Research Centre of the European Commission, more than “80 percent of all product-related environmental impacts are determined during the design phase of a product.”

One company that applies AI at the start of the design process is Digimind, based in Berlin. As CEO Katharina Eissing told *Packaging Europe* last year, Digimind's AI-aided platform lets package designers quickly assess the outcome of changes they make to designs. In one case, Digimind reduced the weight of a company's 1.5-liter plastic bottles by 13.7 percent, a seemingly small improvement that becomes more impressive when you consider that the company produces 1 billion of these bottles every year.

That's still just a drop in the polyethylene terephthalate bucket: The world produced an estimated 583 billion PET bottles last year, according to Statista. To truly address our global garbage problem, our consumption patterns must change—canteens instead of single-use plastic bottles, compostable paper boxes instead of plastic clamshell containers, reusable shopping bags instead of “disposable” plastic ones. And engineers involved in product design need to develop packaging free of PET, polystyrene, and polycarbonate, which break down into tiny particles called microplastics that researchers are now finding in human blood and feces.

As much as we may hope that AI can solve our problems for us, that's wishful thinking. Human ingenuity got us into this mess, and humans will have to regulate, legislate, and otherwise incentivize the private sector to get us out of it. ■

PORTRAIT BY SERGIO ALBIAC; AMP ROBOTICS

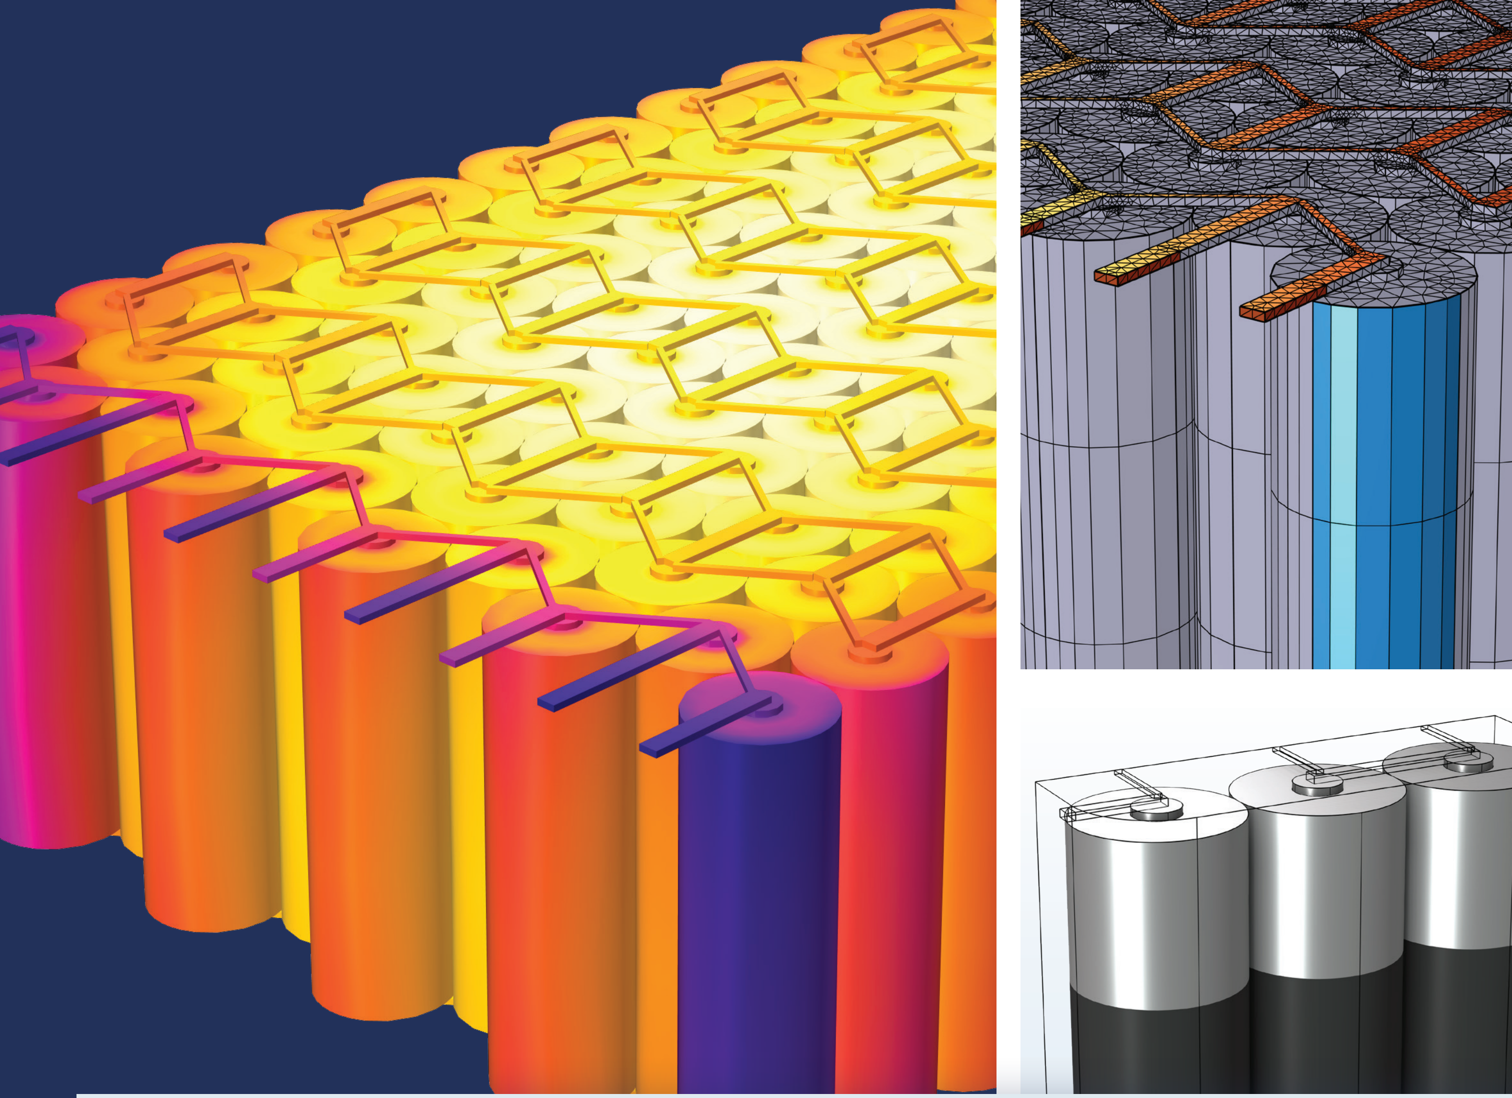

# CONTRIBUTORS

## ● MICHAEL J. BIERCUK

Biercuk is the CEO and founder of Q-CTRL, an Australian startup intent on making quantum computing useful. He is also the director of the Quantum Control Laboratory at the University of Sydney. His coauthor, Thomas M. Stace, a professor at the University of Queensland, is principal quantum control engineer at Q-CTRL. On page 28, they write about the many challenges of quantum error correction, a technique needed to build general-purpose quantum computers.

## ● JASON CALAIARO

Calaiaro is director of hardware for Amp Robotics. He describes the company's AI-based systems for sorting recycling on page 22. Before joining Amp, he founded Marble, now part of Caterpillar, where he pioneered robots for last-mile delivery. He also developed aerial transportation drones at Matternet and served as director of propulsion at Astrobotic Technology, which plans to be the first private company to land on the moon.

## ● ASHOK JHUNJHUNWALA

Jhunjhunwala is an institute professor at the Indian Institute of Technology Madras (IITM), in Chennai, and an IEEE Fellow. In this issue, he and his colleagues Kaushal Kumar Jha and Anson Sando describe a pilot project at the IITM Research Park aimed at accelerating India's use of solar and wind power and energy storage [p. 40]. The same approach could be applied to the country's 40,000 commercial complexes, as well as to its industrial and residential complexes. The research park "has taken the first step to move India towards 100 percent renewable energy," Jhunjhunwala says.

## ● MITCHELL LAZARUS

Lazarus, now retired, has earned his living as an electrical engineer, psychology professor, education reformer, educational-TV developer, freelance writer, and telecommunications attorney. On page 34 of this issue, he writes about conflicts that arise when new radio-based services in overcrowded bands threaten incumbent services with interference. Having represented parties in dozens of such legal disputes, he's had a front-row seat for viewing how these controversies play out.

**IEEE Spectrum**

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

**MANAGING DIRECTOR, PUBLICATIONS** Steven Heffner

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

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

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

**REPRINT SALES** +1 212 221 9595, ext. 319

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

**IEEE 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. 7. The editorial content of IEEE Spectrum magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, IEEE Spectrum, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS: IEEE Members: $21.40 included in dues. Libraries/institutions: $399. POSTMASTER: Please send address changes to IEEE Spectrum, % Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at 120 Donnelley Dr., Glasgow, KY 42141-1060, U.S.A. IEEE Spectrum circulation is audited by BPA Worldwide. IEEE Spectrum is a member of the Association of Business Information & Media Companies, the Association of Magazine Media, and Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit https://www.ieee.org/about/corporate/governance/p9-26.html.

# News

London's Mission Zero Technologies has developed a carbon dioxide sequestration technology, storing atmospheric carbon in the dominant rock (peridotite) of the upper part of the Earth's mantle.

GEOENGINEERING

## XPrize Competitors Capture Carbon >

### $100 million at stake in $CO_2$-removal face-off

BY PRACHI PATEL

**Stretching across the** northern coasts of Oman and the United Arab Emirates loom the vast jagged peaks of the Al Hajar mountains. The craggy outcrops are made mostly of a rock called peridotite, which absorbs carbon dioxide from the air and turns it into solid minerals. The mountains could store trillions of tonnes of human-made $CO_2$ emissions, but the natural carbon-mineralization process works at a glacial pace.

London startup 44.01 has found a way to speed it up. For this endeavor, 44.01 is teaming up with another London startup, Mission Zero Technologies, which has developed an energy-efficient method to capture $CO_2$ from air. Called Project Hajar, it plans to pull 1,000 tonnes of $CO_2$ per year from air at a demonstration facility in Oman, injecting some 3 to 4 tonnes per day into the peridotite rocks. A 120-tonne-capacity pilot plant will come on line in the first half of 2023.

This clear, ambitious vision made Project Hajar one of 15 winners of a milestone US $1 million award announced by the ongoing XPrize Carbon Removal competition in late April. Funded by Elon Musk, this XPrize has the largest purse yet—$100 million—for methods to pull $CO_2$ from air and lock it away. The 15 teams, selected from over 1,100, had to demonstrate a viable approach along with scale-up plans and cost estimates.

Carbon removal is not to be confused with carbon capture at smokestacks. Pulling $CO_2$ from air, where it's present at a very low concentration, is far more complex and costly. Yet in an April report,

Canada's Carbin Minerals is testing a $CO_2$-to-mineral-rock technology that could unlock the potential for gigatonne-scale atmospheric carbon removal and permanent storage.

CARBIN MINERALS

the U.N. Intergovernmental Panel on Climate Change says that carbon removal will be "unavoidable" to keep the planet from crossing the life-disrupting global-warming threshold of 1.5 °C above preindustrial levels. The decarbonizing method is gaining popularity, with high-profile investors including Musk, Bill Gates, and Google's parent company, Alphabet, pouring millions into promising solutions. The Biden administration also recently announced a $3.5 billion program for large-scale carbon removal.

The 15 award-winning approaches include direct air capture (DAC) using chemicals, turning farm waste into charcoal and burying it, growing algae or kelp, and tweaking ocean pH to boost its natural capacity to soak up $CO_2$. The $50 million grand prize, to be awarded in 2025, is up for grabs for any team that can prove its technique will work at a scale of at least 1,000 tonnes per year.

That immense scale, as well as what happens to the $CO_2$, will decide whether an approach can make a dent in the world's nearly 36 billion tonnes of annual carbon emissions, says Gaurav Sant, director of the Institute for Carbon Management at the University of California, Los Angeles, who has two entries in the competition (SeaChange and BeyonDAC), neither of which was among the 15 milestone awardees. Any meaningful approach needs to convert the gas into something stable and not just bury it in the ground from where it could leak.

"Prizes are important because they provoke optimism," Sant says. "We require the development of a portfolio of solutions while we take lots of shots on goal. At the same time we need to be robust and thoughtful, both about the technology development but also the eventual fate of $CO_2$."

Not all 15 winners have the same potential to meaningfully reduce emissions. Five projects, for instance, rely on land-based techniques like biomass-based power generation, farming algae, planting trees, or changing soil composition by adding charcoal from waste, says Christopher Jones, a chemical and biomolecular engineer who studies carbon capture at Georgia Tech. These approaches are low cost, at under $100 per ton of captured $CO_2$, "but there's only so much land change you could make to capture a significant amount," he says. "We need to capture 10 gigatons per year for negative emissions by 2060. Land and biomass approaches only scale to a few gigatons."

Of known carbon-removal techniques, two hold the most promise, Jones says, citing a recent National Academies report. One is DAC and the other is carbon mineralization. "Nothing prevents us from scaling these up to the 10-gigatons-per-year scale that is needed, aside from a commitment, coordination, and cooperation."

Direct air capture has already taken off, with about 20 projects already underway around the world. Most rely on large fans to suck $CO_2$ from air using liquid or solid materials—which are not cheap—and then heating the mixture by burning natural gas to remove $CO_2$and regenerate the adsorbing material. The downside of DAC is high fossil-fuel energy use and cost.

Sustaera in Cary, N.C., one of six milestone winners pursuing DAC, has found a way around this carbon conundrum. Chief technology officer Raghubir Gupta worked for two decades on carbon capture at power plants and industrial plants. "One big thing we have that not many others have is practical experience of scaling up the technology to 1,000-tonnes-per-day carbon dioxide capture," he says. "With that background, when we looked at $CO_2$ removal from air we thought the two things that are most important to really make a difference are cost and scale. It's not the efficiency of the process."

Sustaera uses cheap sodium carbonate to adsorb $CO_2$. It coats the material on a high-surface-area ceramic scaffold used in catalytic converters. The high surface area increases access to the sorbent and increases the $CO_2$ adsorption rate significantly, Gupta says.

Instead of burning fossil fuels for heat, Sustaera uses Joule heating, in which passing an electric current through a conductor produces heat to separate the $CO_2$ and regenerate the sorbent. Along with the sodium carbonate, the ceramic support contains a conducting material such as carbon nanotubes. Electricity, which can be renewable, locally heats the sorbent and triggers $CO_2$ release. At full scale, Sustaera's system should be able to capture more than 3,000 tonnes per day of $CO_2$ at under $100 per tonne. For now, says Gupta, a 1-tonne-per-day facility being built at the company's R&D site in Research Triangle Park, N.C., should be ready by the end of next year.

While Sustaera focuses on making DAC cheap, Project Hajar's promise lies in marrying DAC with permanent carbon storage via mineralization. First, project partner Mission Zero Technologies uses solvents to capture $CO_2$ from air blown through a tower. Then an electrochemical cell separates the $CO_2$. The process takes a third of the energy of conventional thermal separation. "It works fully with existing materials and chemicals and off-the-shelf equipment," says Mission Zero's cofounder Shiladitya Ghosh. Both the cooling tower and the electrochemical-cell technologies are ubiquitous around the world, "so the manufacturing systems are established and available."

Then, startup 44.01 mixes the $CO_2$ with water and injects it via engineered boreholes into the peridotite rocks to form carbonate minerals in less than a year. "We accelerate the reaction by creating physical and chemical characteristics to catalyze it, such as pressure, temperature, and alkalinity balance in the subsurface," says the company's cofounder Karan Khimji.

Like its partner, 44.01 also uses off-the-shelf equipment from the oil and gas sector. "I find beauty in that," Khimji says. "We're repurposing the same resources as the oil and gas sector to reverse the problem that they have contributed to." Renewable electricity on-site will power both carbon removal and mineralization at the future Oman demonstration facility. Another competitive advantage is permanence: "$CO_2$ is eliminated from existence; it doesn't remain in gaseous form in the subsurface," he says.

For sheer scale, nothing could beat the oceans as a carbon sink, says UCLA's Sant. Three of the XPrize milestone winners have ocean-based carbon-removal platforms. But the caveat for large-scale impact is to stabilize the $CO_2$ in ocean water, not in a geological formation, he says. The way to do that is to enhance the ocean's natural uptake of $CO_2$.

Here, Planetary Technologies from Nova Scotia, Canada, might have the most interesting approach. Rising carbon levels are making the world's oceans acidic. The company purifies mine waste to make a mild antacid to restore ocean-water pH, which should help it pull more $CO_2$ from air while reducing damage from acidification. It says that its mine-waste purification technique also produces hydrogen for energy and metals for batteries. In this way it tackles several different issues at once: carbon removal, green hydrogen production, mine-waste cleanup, and ocean restoration.

The big winner for the XPrize Carbon Removal will be announced on Earth Day 2025. But of course, there is no one winning solution, says Jones. All carbon-removal technologies have pluses and minuses and bear the risk of unintended consequences. "The problem is big enough that you need a dozen different technologies to contribute a little," he says. ■

JOURNAL WATCH

## Touchless Touching in VR

Virtual-reality headsets allow people to immerse themselves in environments completely divorced from the real world. But some researchers are interested in expanding connections to these fantasy worlds beyond just visual and acoustic sensations. For example, many groups have been searching for ways to incorporate tactile sensation into VR.

One difficulty that has heretofore impeded the simulation of physical touch is the bulky hardware that must be worn on the hands to emulate the sensations. This hardware is especially cumbersome if the aim is to create sensations with different strengths and directions of force. Machinery would need to nudge the user's hands from multiple angles.

But Taha Moriyama, a researcher at the University of Electro-Communications, in Chōfu, Japan, says he has gotten a grasp on virtual touching—without all the bulk that would make it difficult for the user to forget the world outside the VR headset. Moriyama reports that he has developed a new approach to VR haptics that sidesteps the need for hardware on a user's hands. Instead, haptic sensations are applied to the person's forearm.

In his latest research with colleague Hiroyuki Kajimoto, the two describe this novel approach in a study that was published in *IEEE Transactions on Haptics* in January.

The haptic system they invented can be 3D printed and weighs just 250 grams. It includes an external sensor camera that tracks the user's finger movements. Haptic sensations are then applied to the top, bottom, or sides of the forearm to deliver feedback that matches the movements of the fingers. For example, moving a finger from left to right along the surface of a VR object would trigger the device to apply a left-right surface-brushing sensation to the nerves in the forearm.

For their study, Moriyama and Kajimoto recruited 11 volunteers and assessed their comfort levels with the device as they each completed the simple task of grasping a VR object and moving it to a predetermined position.

"We were surprised that even with this new haptic presentation method, we were able to obtain a high comfort level without any training time," says Moriyama. He adds, "In future work, we intend to design a device, based on our proposed device, that can [deliver sensations mimicking] vibration and [changes in] temperature."

**— By Michelle Hampson**

A Sikorsky S-92 helicopter [top] hovers in the catch zone waiting to halt the descent of the rocket's first stage, while a two-stage parachute [bottom] slows the rocket booster enough to make the midair grab possible.

ROCKET LAB

AEROSPACE

# Rocket Lab Catches Rocket Booster in Midair

## Successful booster recoveries can dramatically cut the cost of space launches

BY NED POTTER

**The longest journey** begins with s single step—but when you're in the space business, each step can be costly. Take, for example, the Electron rocket made by Rocket Lab USA, a company with two launch pads on the New Zealand coast and another awaiting use in Virginia. Earth's gravity is so stubborn that, by necessity, two-thirds of the 18-meter-tall rocket is its first stage—a rocket segment that has historically ended up as trash on the ocean floor after spending a little over 2 minutes in flight. On many missions, the first stage burns out after propelling a rocket and its payload for a little more than the first 70 kilometers of altitude after liftoff. Then it drops off, following a long arc that sends it crashing into the ocean, about 280 km downrange.

Making those boosters reusable—saving them from a saltwater grave, and therefore saving a lot of money—has been a goal of aerospace engineers since the early space age. Elon Musk's SpaceX has famously been landing its Falcon 9 boosters on drone ships off the Florida coast—a stunning maneuver, but very hard to pull off.

Rocket Lab has imagined another way: Instead of letting the spent first stage crash in the Pacific, it plans to have specially equipped helicopters catch its rockets in midair as they descend by parachute.

On 2 May, Rocket Lab made its first attempt to pull off the maneuver. It was only partially successful, which shouldn't be a surprise, because catching stuff in mid-air as it plummets from high up in the atmosphere is a challenging feat.

When the booster separates from the rest of the rocket, it falls tail-first at supersonic speeds approaching 8,300 kilometers per hour. Temperatures on the shield's exterior reach 2,400 °C as it's buffeted by the air around it.

At an altitude of 13 km, a small drogue parachute is deployed from the top end of the rocket stage, followed by a main chute when it falls to 6 km, less than a minute later. The parachutes slow the rocket so that it is soon descending at only about 36 km/h.

But even that would make for a hard splashdown—which is why the helicopter, a Sikorsky S-92, hovers over the landing zone, trailing a grappling hook on a long cable. The helicopter flies over the descending rocket and snags the parachute cables with the aim of keeping the rocket from getting wet. The chopper's job is to then lower it onto a waiting ship or carry it back to land.

It goes without saying that making the catch is easier said than done. The midair maneuver has to be set up just right in order to be successful.

"You have to position the helicopter in exactly the right spot; you have to know exactly where the stage is going to be coming down; [and] you have to be able to slow it enough," says Morgan Bailey of Rocket Lab.

So, how did the first midair recovery attempt go?

Rocket Lab's launch didn't have a storybook ending, though its attempt to catch the first stage was no less spectacular than was predicted. Launch and ascent were virtually flawless; 34 satellites were released in orbit. And roughly 14 minutes after liftoff, video from the helicopter showed that it had just managed to catch the descending booster. But because the aircraft was attempting to snag a 12-meter-long tin can whizzing by at speeds like those of a car in city traffic, Rocket Lab's catch faced a complication that the company will need a few more tries to master.

Rocket Lab says that although the helicopter caught the booster, the pilot experienced a variance in the load characteristics that he hadn't experienced during testing. So, following established safety protocols, he released it. The rocket stage continued by parachute to the water below, where a recovery ship hauled it on board. Still, Rocket Lab is undaunted. The company says it will keep at it until midair recoveries become routine. That goes along with the mantra adopted by company CEO Peter Beck: "Launch, catch, repeat."

"Epic day," said Beck, regarding the Electron rocket launch. "The difficulty in capturing a stage is pretty extreme."

Rocket Lab's Bailey, speaking before the trial, acknowledged that success was not a foregone conclusion. "We've practiced and practiced and practiced all of the individual puzzle pieces, and now it's [a matter of] putting them together," she told *IEEE Spectrum* at the time.

Despite there being no guarantee that Rocket Lab would pull off the ambitious maneuver, people in the space business gave it credence, because Rocket Lab had already established a niche for itself as a viable space company. It had previously launched 25 Electron rockets, carrying a total of 112 satellites into orbit—most of them so-called smallsats that are relatively inexpensive to fly.

Meanwhile—let's not lose sight of the prime mission—the second stage of the rocket reached orbit about 10 minutes after launch, the company reports. ■

---

BATTERIES

# Manganese: The Secret Behind Truly Mass-Market EVs? › This abundant transition metal could electrify the automobile for the global mainstream

BY LAWRENCE ULRICH

**Most automakers are** dying to sell you—and the world—an electric car. But they're up against a challenge that threatens to make their engineers' innovations come to naught: dauntingly tight supplies of batteries and of ethically sourced raw materials for producing them.

Tesla and Volkswagen are among the automakers who see manganese—element No. 25 on the periodic table, situated between chromium and iron—as the latest, alluringly plentiful metal that might make batteries, and therefore electric vehicles, affordable enough for mainstream buyers.

That's despite the dispiriting history of the first (and to date, the only) EV to use a high-manganese battery: the original Nissan Leaf, back in 2010. But with the industry needing all the batteries it can get, improved high-manganese batteries could carve out a niche, perhaps as a midpriced option between lithium-iron phosphate versions and the primo nickel-rich batteries that would power top luxury and performance models.

Elon Musk made waves this March at the opening ceremony of Tesla's Gigafactory Berlin-Brandenburg when asked his opinion on graphene in battery cells: "I think there's an interesting potential for manganese," he countered.

Regarding raw minerals, Musk underlined the ongoing industry flight from cobalt and now nickel: "We need tens, maybe hundreds of millions of tons, ultimately. So, the materials used to produce these batteries need to be common materials, or you can't scale," he said.

In March 2021, VW unveiled a versatile "unified cell"—part of its plan to establish gigafactories that would deliver a total of 240 gigawatt-hours of capacity by 2030. The unified cell is compatible with multiple chemistries and comes in a standardized prismatic design. VW CEO Herbert Diess said then that about 80 percent of VW's new prismatic batteries would spurn pricey nickel and cobalt in favor of cheaper, more-plentiful cathode materials—among them, potentially, manganese.

So, why this endless mixing-and-matching of formats and cathodes? And why manganese? It all hinges on what Musk and other experts cite as the looming limiting factor in accelerating the EV revolution: the lagging rate of both battery production and the mining and processing of their raw materials.

In Berlin, Musk suggested the world will need 300 terawatt-hours of EV battery storage available in order to realize a full transition from fossil-fueled cars. That's more than a dozen times what Tesla projects it can produce by 2030, even with its own massive expansion of capacity. Nickel-rich batteries alone won't get us there, despite unprecedented energy density and performance. Other materials will be required; sourcing them with ethical, diverse, uninterrupted pipelines will be critical to successfully scaling production. This is true even if chemistries featuring manganese or lithium-iron phosphate—the flavor of the moment for EVs—yield batteries that are the result of some compromises.

A Tesla electric car gets a charge at an EnBW fast-charging park in Germany.

Manganese is abundant, safe, and stable—qualities that make it a promising candidate for the cathode material as EV battery production is ramped up. But it has never approached the energy density or life cycle of nickel-rich batteries, cautions Venkat Srinivasan, director of the Collaborative Center for Energy Storage Science at Argonne National Laboratory. Buyers of early Nissan Leafs might concur: Nissan, with no suppliers willing or able to deliver batteries at scale back in 2011, was forced to build its own lithium manganese oxide batteries with a molecular jungle-gym-like "spinel" design. Those energy-poor packs brought just 24 kilowatt-hours of storage and a 135-kilometer (84-mile) driving range. Even that piddling storage and range rapidly degraded, especially in the southwestern United States and other searing climates, leaving customers howling. (It didn't help that Nissan eschewed a thermal-management system for the battery.) A "lizard" battery in 2014 with a modified manganese chemistry still had a capacity of 24 kWh, and also suffered short life spans.

Still, manganese remains in contention as an EV battery metal because its high-performing counterpart, cobalt, is not only pricey but also comes mainly from the Democratic Republic of the Congo, which is linked to child labor in mines and other human rights abuses. Low-cobalt batteries have been the industry's response to those supply-chain issues.

The next popular cathode mineral has been nickel, with a more diverse supply than Congolese cobalt but hardly immune from geopolitical concerns. Global nickel stockpiles were already dwindling before Russia's invasion of Ukraine in February. Investors and traders got antsy over potential bans or interruptions of metals from Russia, which produces about 17 percent of the world's high-purity nickel. In March, nickel prices doubled virtually overnight, briefly topping US $100,000 per tonne for the first time, spurring the London Metal Exchange to suspend trading during the wild run-up.

For all these reasons—commodity prices, politics, ethics, security, shortages, long-term strategy, and hedging of bets—the industry is embarking on a diversification strategy, a smorgasbord of solutions. Or at least until some future Nobel winner comes up with something to replace lithium-ion entirely.

"Everyone is thinking about substitutions for nickel and cobalt and how to recycle these things," says Srinivasan.

Pouch-style Ultium cells from General Motors and LG Energy Solution—which I recently tested for the first time in the GMC Hummer EV—use a nickel cobalt manganese aluminum chemistry that reduces cobalt content by about 70 percent. With 200 kWh in a double-stacked cell sandwich—twice the size of Tesla's biggest battery—the reborn Hummer combines a 529-km (329-mile) range with trimotor propulsion, 1,000 horsepower, and a 3.0-second explosion to 60 miles per hour in its WTF ("Watts to Freedom") mode. That battery, by far the largest ever shoehorned into an EV, also contributes 1,326 kilograms to the Hummer's gargantuan 4,111-kg curb weight. (With GM gearing up mass production in Detroit, the Hummer will create massive battery storage demand all on its own.)

GM's cells use only small amounts of manganese to stabilize structures, not as a main cathode material.

For the fickle automaker, for which even nickel is on the outs, the switch is on to lithium-iron phosphate (LFP) chemistries. That battery chemistry—invented in the 1990s and until recently viewed as yesterday's news—requires no nickel or cobalt, just abundant iron and phosphate. Musk has confirmed a "long-term switch" to LFP for entry-level cars (including the Tesla Model 3) or energy storage.

High-manganese batteries being eyeballed by Tesla and VW would also use less nickel and zero cobalt. They appear affordable: According to analysts at Roskill Information Services, in London, a lithium nickel manganese oxide chemistry could reduce cathode costs by 47 percent per kilowatt-hour relative to nickel-rich designs. That has VW mulling manganese as a potential fit for mainstream models, LFP for bottom-rung vehicles or markets, and bespoke high-performance packs for the likes of Porsche, Audi, Bentley, or Lamborghini.

"I can see the logic, where if you can get it to a reasonable energy density, manganese becomes this in-between thing," Srinivasan says. Automakers might offset manganese's lower cathode costs with slightly enlarged batteries, to bring range closer to par with nickel-rich designs.

Back in 2020, at Tesla's Battery Day, Musk expressed optimism about the mineral: "It is relatively straightforward to do a cathode that's two-thirds nickel and one-third manganese, which will allow us to make 50 percent more cell volume with the same amount of nickel."

With Musk still struggling to bring his large-format 4680 cylindrical cell to market—now well behind schedule—experts warn that the technical challenges aren't so straightforward. High-manganese batteries have yet to demonstrate commercial viability.

But the epic scale of the challenge has automakers and battery makers working the labs and scouring the globe for materials as common as dirt, not precious as gold. ■

BERND THISSEN/PICTURE ALLIANCE/GETTY IMAGES

SEMICONDUCTORS

# Single-Chip Processors Have Reached Their Limits > Chiplets seem to be the future, but interconnects remain a battleground

BY MATTHEW S. SMITH

**Apple once again** surprised enthusiasts and analysts with its announcement of the M1 Ultra, a variant of the M1 Max that effectively fuses two chips into one. The result is a dual-chip design viewed by software as a single piece of silicon. Nvidia delivered similar news in March at its GPU Technology Conference. CEO Jensen Huang announced that the company will fuse two of the company's new Grace CPU processors into a single "Superchip."

These announcements target different markets. Apple has its sights set on the consumer and professional workstation world, while Nvidia intends to compete in high-performance computing. Yet the divergence in purpose only underscores the array of challenges rapidly bringing the era of monolithic chip design to an end.

Multichip design isn't new, but the idea has surged in popularity in the past five years. Advanced Micro Devices, Apple, Intel, and Nvidia have all dabbled to varying degrees. AMD has pursued chiplet design with its Epyc and Ryzen processors. Intel plans to follow suit with Sapphire Rapids, an upcoming architecture for the server market built on the use of chiplets it calls "tiles." Now, Apple and Nvidia have joined the party—though with designs focused on significantly different use cases.

The shift toward multichip design is driven by the challenges inherent in modern chip manufacturing. Miniaturization of transistors has slowed, yet growth in transistor counts in leading-edge designs shows no sign of abating.

Apple's M1 Ultra has 114 billion transistors and a die area (or fabrication area) of roughly 860 square millimeters. (An official figure for the M1 Ultra is unavailable, but a single M1 Max chip has a die area of 432 mm$^2$.) The transistor count of Nvidia's Grace CPU is still under wraps, but the Hopper H100 GPU announced alongside the Grace CPU includes 80 billion transistors. For perspective, AMD's 64-core Epyc Rome processor, released in 2019, has 39.5 billion transistors.

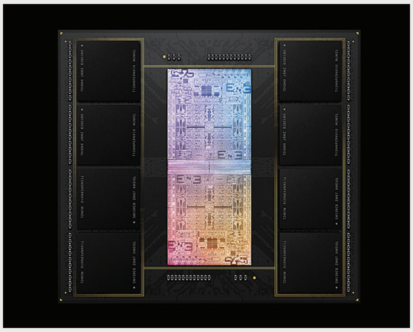
Apple's M1 Ultra is a dual-chip design that software sees as a single piece of silicon.

APPLE

"Multichip module packaging has enabled the chipmakers to give [billion-plus-transistor setups] better power efficiency and performance [in comparison with] monolithic designs, as the die size for chips becomes larger and wafer-yield issues become more prominent," Akshara Bassi, an analyst at Counterpoint Research, said in an email. Aside from Cerebras, a startup attempting to build chips that span the entirety of a silicon wafer, the chip industry seems to be in agreement that monolithic design is becoming more trouble than it's worth.

This shift towards chiplets has occurred in step with support from chip fabs. Taiwan Semiconductor Manufacturing Co. (TSMC) is an early adopter of the trend, offering a suite of advanced packaging called 3DFabric. Technologies that fall under the umbrella of 3DFabric are used by AMD in some Epyc and Ryzen processor designs and are almost certainly used by Apple for M1 Ultra. (Apple has not confirmed this, but the M1 Ultra is produced by TSMC.) Intel has its own packaging technologies, such as EMIB and Foveros. Though originally meant for Intel's exclusive use, the company's chip-manufacturing technology is becoming relevant to the broader industry as Intel Foundry Services opens its doors.

"The ecosystem around the foundational semiconductor design, manufacturing, and packaging [of chiplets] has progressed to the point of supporting the design nodes to economically and reliably produce chiplet-based solutions," Mark Nossokoff, a senior analyst at Hyperion Research, said in an email. "The software design tools to seamlessly integrate the chiplets' functionality have also matured [significantly]," says Nossokoff.

Chiplets are here to stay, but for the moment, it's a world of silos. AMD, Apple, Intel, and Nvidia are using their own interconnect designs meant for specific packaging technologies.

Universal Chiplet Interconnect Express (UCIe) hopes to bring the industry together. Announced on 2 March 2022, this open standard offers a "standard" 2D package that targets "cost-effective performance" and an "advanced" package that targets leading-edge designs. UCIe also supports off-package connection through PCIe and CXL, making possible the connection of multiple chips across multiple machines in a high-performance compute environment.

UCIe is a start, but the degree to which the standard will be embraced by the industry remains to be seen. "The founding members of initial UCIe promoters represent an impressive list of contributors across a broad range of technology design and manufacturing areas, including the HPC ecosystem," said Nossokoff, "but a number of major organizations have not yet joined, including Apple, AWS, Broadcom, IBM, and Nvidia."

Nvidia has debuted its own NVLink-C2C interconnect for custom silicon integration, making it a potential competitor for UCIe.

Platforms like UCIe and NVLink-C2C might determine some rules of the game, but they're unlikely to change the name of the game being played. Apple's M1 Ultra could be considered the canary in the coal mine, indicating that multichip design is coming to a home computer near you. ■

# Megatruck Runs on the Lightest Gas

By Willie D. Jones

Big things are happening in the world of hydrogen-powered vehicles. One of the latest monumental happenings is the debut of Anglo American's 510-tonne hydrogen-powered mining truck. The behemoth will replace an entire fleet of 40 diesel-powered trucks that haul ore away from a South African platinum mine. Together, those trucks consume about one million liters of diesel fuel each year. The new truck, whose power plant features eight 100-kilowatt hydrogen fuel cells and a 1.2-megawatt battery pack, is just the first earth-moving step in Anglo American's NuGen project, aimed at replacing its global fleet of 400 diesel mining trucks with hydrogen-powered versions. According to the company's estimates, the switch will be the equivalent of taking half a million diesel-fueled passenger cars off the road.

PHOTOGRAPH BY WALDO SWIEGERS/ BLOOMBERG/GETTY IMAGES

# Hands On

Radioactive minerals can be identified in a surprising number of places, including old ceramic glazes.

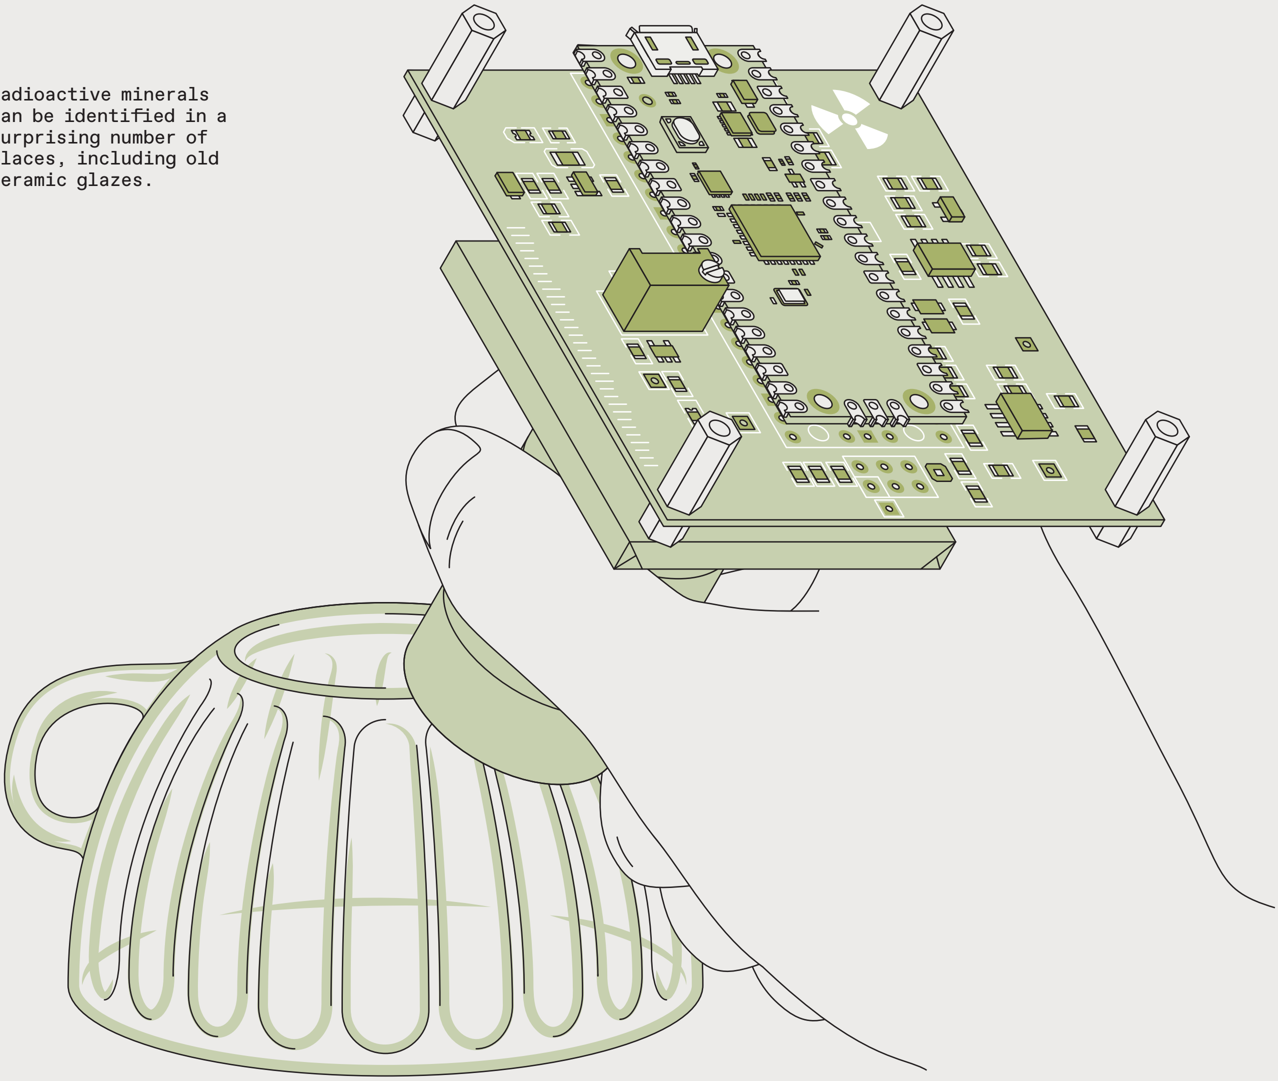

## DIY Gamma-Ray Spectroscopy › But beware a glitch in the popular Pi Pico microcontroller

BY MATTHIAS ROSEZKY

**The global semiconductor** shortage has made life tough for anyone using microcontrollers, with lead times now sometimes quoted in years. But there has been one bright spot: the US $4 Pi Pico, a microcontroller based on the new RP2040 chip. Not only does the RP2040 have lots of compute power, it hasn't suffered the kind of shortages afflicting other chips. So when I decided to build a cheap DIY scintillating gamma spectrometer, it was the natural choice—although I didn't

*Illustrations by James Provost*

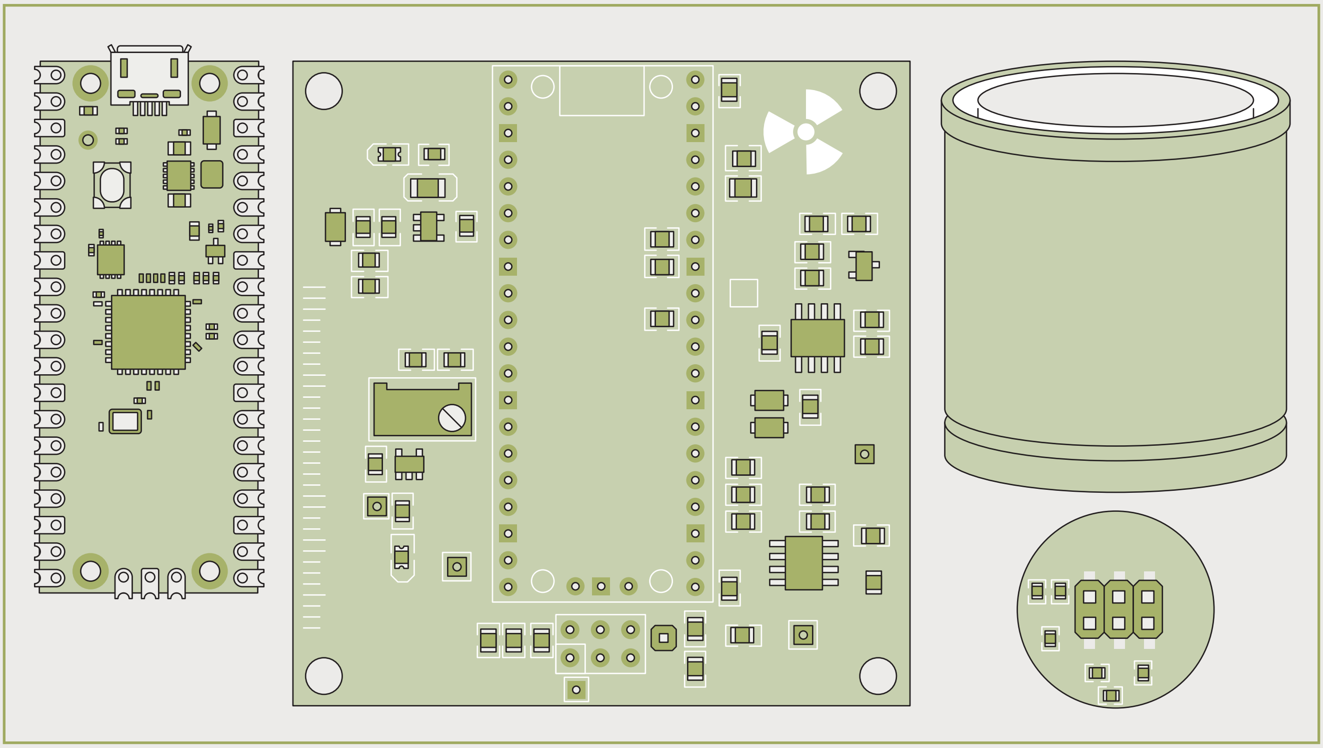

A Raspberry Pi Pico [left] provides both compute power and the gamma-ray spectrometer's analog-to-digital converter. A board [middle] provides power and an amplifier connected to the silicon photomultiplier carrier board [bottom right] and scintillating crystal [top right], which reacts to gamma rays.

realize I'd find myself navigating around teething problems of the sort that often affect a first-generation integrated circuit.

My interest in gamma-ray spectroscopy came from my physics studies. I find it fascinating that you can get so much information out of a single device. A gamma-ray spectrometer can be used like a Geiger counter with much better sensitivity, but unlike with a Geiger counter, it lets you identify the exact composition of any gamma-emitting radioisotopes down to the picogram (or less). I started thinking about creating my own gamma-ray spectrometer when I saw the high price of even the cheapest commercially made devices. I wanted to see if I could make it easy and affordable to build a spectrometer.

The first step was to choose the scintillator at the heart of the spectrometer. In a nutshell, a scintillator measures both the energy and intensity of a flux of gamma rays, thanks to a transparent crystal. A gamma ray produces a free electron in the crystal, and this electron's energy is proportional to the gamma ray's. As the electron moves through the crystal, it excites atoms. The atoms, in turn, emit visible photons, with the total number of photons emitted proportional to the energy of the exciting electron. Thus, by counting the number of photons, you can gauge the energy of the original gamma ray. Counting how many gamma rays you detect over time gives you the radiation's intensity, and looking at the energies of the gamma rays gives you a spectral fingerprint of a radioisotope.

The photon signal must be amplified to be detectable. Historically, this was done using a photomultiplier vacuum tube, but silicon photomultipliers (SiPMs) have become more common, and for my project they have a number of advantages, particularly in eliminating the need for a high-voltage power supply.

You can buy various used scintillator crystals on eBay fairly cheaply: I purchased a small sodium iodide crystal, 18 millimeters in diameter and 30 mm long, for about US $40. It came with a photomultiplier tube, which I removed and replaced with my SiPM, wrapping the assembly in black tape to prevent external light from leaking in and triggering the sensor.

The scintillator and SiPM plug into a carrier board, which has a DC/DC boost converter to convert 5 volts into the 29.3 V the SiPM needs. The carrier board also hosts the Pico microcontroller along with some other supporting circuitry, including an amplifier that increases the

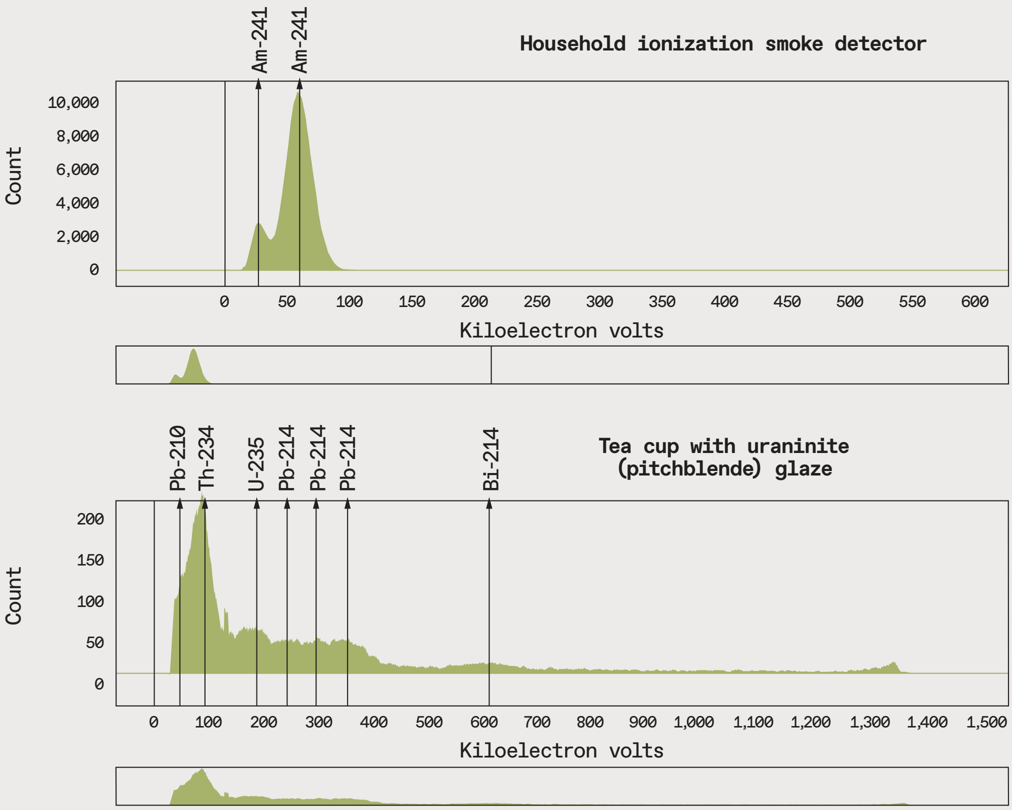

Radioactive minerals are more common than many people think: two sample spectra captured with the DIY detector and the isotopes responsible for their signatures. The boxes beneath each spectrum show the zoomed-out readings across all channels

output voltage of the SiPM to a level that the Pico's built-in analog-to-digital converter (ADC) can detect.

The ADC in the Pico's RP2040 chip is a critical component, and on paper it looks very good.

It has 12-bit resolution and can take measurements between 0 and 3.3 V at a rate of 500 kilosamples per second. But there's a flaw lurking in the RP2040's ADC.

I didn't realize it existed until I started taking test spectra, writing software for the Pi that breaks up the ADC's readings into 4,096 channels and counts the number of events in each channel over time. I noticed that one channel kept reporting very high count values, creating a thin spike in my spectra. Puzzled, I took a 4-hour background radiation measurement and discovered there were four problematic channels where the signal spiked unrealistically.

I started searching for what could be causing this and discovered I was not the first to run into problems with the ADC. A great website by Mark Omo—an EE who took it upon himself to investigate the problem—provides a detailed analysis, but in summary the issue is this: Ideally, an ADC chops the voltage range it can measure into an identically sized sequence of steps, producing a linear relationship between the input voltage and the numeric measurements it outputs. Of course, no ADC has a perfectly linear response across its measurement range, but the RP2040 has four spots where input voltages produce a wildly nonlinear response. This was the source of the mystery spikes in my spectra.

Until the RP2040 is revised to fix this glitch, there's not much you can do about it directly. Fortunately, with 4,096 channels, I could afford to employ the simplest software fix—just throwing away the measurements in the affected channels—without affecting the quality of the overall spectrum significantly.

Controlling and getting data from the spectrometer can be done via a USB interface (which also provides the power needed to operate it). I wrote software that can accept serial commands to, for example, put the spectrometer into Geiger counter or energy-measurement modes, or upload a histogram of all the measurements taken since the last power-up. You can write your own code to communicate with the spectrometer, or use a Web app I created that also plots spectra. (A link to the Web app, along with full build details and PCB files, is available on GitHub.)

For the future, I hope to make the spectrometer hardware capable of using a wider range of SiPMs and scintillators, so that people can use whatever detectors they can find. I hope you join me in this fascinating hobby! ■

**On paper, the Pi Pico's ADC looks very good. But there's a flaw lurking in it.**

# Careers

DOUGLAS PALMER

In his workshop, venture capitalist Adam Grosser is converting the original drivetrain on his 1974 Jaguar E-Type to an electric one.

## Adam Grosser › This VC is big on transportation alternatives

BY DANIEL P. DERN

**Adam Grosser wants** to make transportation, of people, goods, and energy faster and more efficient for everyone. That's why the chairman and managing partner of the early-stage venture capital fund UP Partners is investing in several mobility projects. They include Beta Technologies' electric vertical-take-off-and-landing (eVTOL) urban transportation aircraft, Quincus's operating system for supply-chain and logistics providers, and Teleo's teleoperation platform for mining, construction, and other heavy equipment.

"Transportation is the underlying fabric of society," Grosser says. "At UP, we invest in key enabling technologies that help move people and goods faster, safer, more efficiently, and sustainably. This can include anything from new kinds of ground, sea-born, air, or space vehicles to production lines, packages, and units of automation."

The mobility sector is ripe for significant improvements, he says. "Arguably just about everything on a car, except for a few safety systems, was invented by 1920—although not necessarily put into widespread practice. But mobility hasn't previously been an investable category."

He credits several factors for this. Faster, smaller, and cheaper additive manufacturing is now available. Also, rapid shifts in battery capacity and electric motor torque have dramatically changed how mobility vehicles and methods are built and work.

The question, he says, is how to pull all these together into something that is safe and more environmentally friendly than what we do today—and which has a viable financial model.

One company that ticks all these boxes for Grosser is Kolors, in Acapulco, Mexico, which aims to transform the bus industry across Latin America by offering a website for riders to reserve a seat on an intercity bus. The company does not own the buses. Instead, it partners with small and medium-size bus lines that own their vehicles to provide a consistent customer experience and offer a single ticketing framework.

Grosser decided to get into the investment business after two decades in senior positions at Apple, Sony Pictures Entertainment, and @Home Network. He spent more than a decade at Foundation Capital, primarily in early-stage ventures. He then moved to large-cap private equity firms. He's been with UP Partners since May 2020.

"Most people who go into tech investing today do that with a fairly clear intention of being an investor," Grosser says. "[By contrast], I would consider myself an inadvertent investor. I've spent decades working to solve meaningful challenges, first as an engineer, then as an entrepreneur, and for the past 21 years as an investor."

What helped him succeed at his investing goals? Mentors.

"I have been lucky enough to have amazing mentors and partners, from my college days through to the present," Grosser says. One was the late Kathryn Gould, who founded Foundation Capital. "She pulled me in and said, 'I think you will be a good investor. Let me teach you.'"

Grosser also credits his diverse engineering experiences with helping him talk knowledgeably about potential new technologies and companies to invest in as well as conduct due diligence. He builds and restores classic cars. He is converting the drivetrain of a 1974 Jaguar E-Type to an electric one.

His advice for would-be investors is to use the knowledge they already have.

"Courses like robotics or thermodynamics that may not have been part of your major can be integrated with whatever you've learned and done in software, hardware, and product design," he says. "Any of this knowledge and experience can help you establish rapport and make more-informed selections." ■

# A Moore's Law—for Bombs

The rising power of destructiveness is, unfortunately, the most impressive metric of modern technology

**The rising number** of components on a microchip is the go-to example of roaring innovation. Intel's first microprocessor, the 4004, released in 1971, had 2,300 transistors; half a century later the highest count surpasses 50 billion, for the Apple M1 Max—an increase of seven orders of magnitude. Most other technical advances have lagged behind: During the entire 20th century, maximum travel speeds rose less than tenfold, from about 100 kilometers per hour for express trains to 900 km/h for cruising jetliners. Skyscrapers got only 2.4 times as tall, from the Singer Building (187 meters) to the Petronas Towers (452 meters).

But there is one accomplishment that, unfortunately, has seen even higher gains since 1945: the destructive power of explosives.

Modern explosives date to the 19th century, with trinitrotoluene (TNT) and dynamite in the 1860s, followed by RDX (Royal Demolition Explosive), patented in 1898. During the Second World War, explosive power rained on European and Japanese cities in the form of mass-scale bombing, and by the war's end, in 1945, the most powerful explosive weapon was the Nazi V-2 rocket. It carried 910 kilograms of amatol—a blend of TNT and ammonium nitrate—and had an explosive energy of about 3.5 gigajoules.

> The increase in explosive power, over 16 years, matches what Moore's Law has accomplished in the 50 years since 1970

And then came an entirely new class of explosives, those exploiting nuclear fission and fusion. The bomb that exploded over Hiroshima on 7 August 1945 released 15 kilotons of TNT (63 terajoules), half of its energy as the blast wave, about a third as thermal radiation. The Nagasaki bomb, dropped two days later, released about 25 kilotons (105 TJ). But these first two bombs were tiny when compared to what came later. The most powerful U.S. hydrogen (or fusion) bomb, tested in 1954, was equivalent to 15 megatons (63 petajoules). This was far surpassed on 30 October 1961, when the Soviet Union tested the RDS-220 bomb above Novaya Zemlya in the Arctic Ocean. Fifty-nine years later, in August 2020, Rosatom (Russia's atomic energy agency) released a 40-minute-long film that claimed that the bomb, nicknamed the *tsar bomba*—the emperor's bomb—had had a yield of 50 megatons.

In this remarkable video, the antiquated analog instrumentation provides a strange contrast with the weapon's immense destructive power. The bomb—hung beneath the belly of a Tu-95 bomber—was dropped by parachute from a height of 10.5 kilometers and detonated 4 km above the ground. The explosion released 210 PJ of energy, three orders of magnitude more than the Nagasaki bomb, creating a mushroom cloud 60-65 km in diameter and a flash visible from nearly 1,000 km away. And soon afterward Nikita Khrushchev, the Soviet premier, claimed that his country had built but not tested a bomb twice as powerful.

The last V-2 attack on London came on 27 March 1945, less than six weeks before the Nazi surrender. By the Novaya Zemlya test in 1961, the maximum explosive energy of weapons had risen by seven orders of magnitude, to more than 200 PJ. That increase, over 16 years, matches what Moore's Law has accomplished in the 50 years since 1970. It is a reminder of the terrible priorities of modern civilization. ■

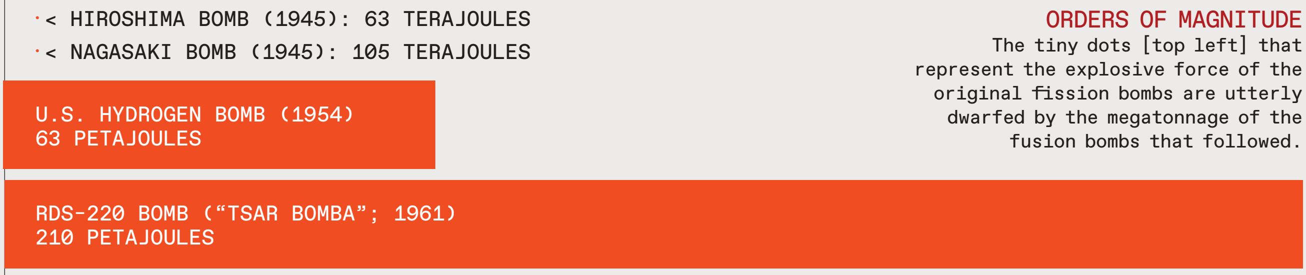

ORDERS OF MAGNITUDE
The tiny dots [top left] that represent the explosive force of the original fission bombs are utterly dwarfed by the megatonnage of the fusion bombs that followed.

CHART SOURCE: NUCLEARWEAPONARCHIVE.ORG

*Portrait by Sergio Albiac*

# The World's Best ROBOTS GUIDE Is Here!

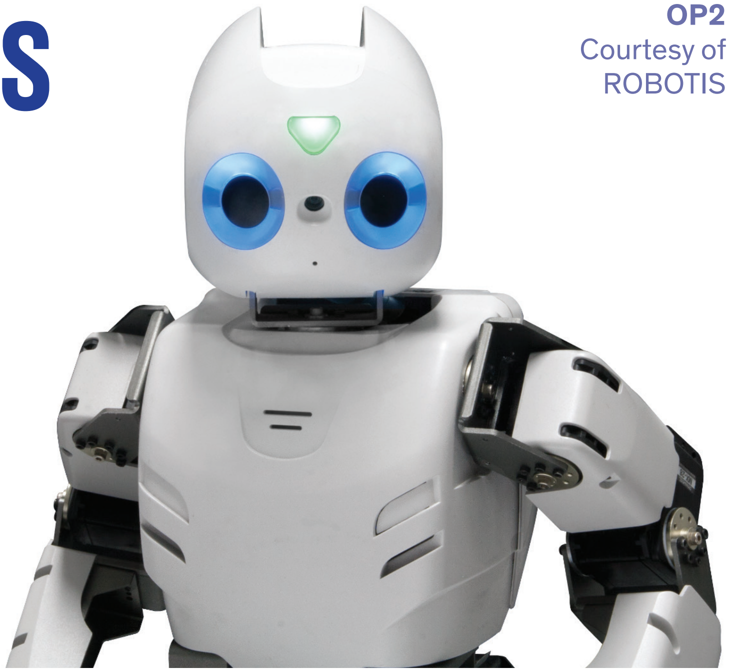

**OP2**
Courtesy of
ROBOTIS

## ROBOTS.IEEE.ORG

***IEEE Spectrum's*** **new ROBOTS site features more than 200 robots from around the world.**

- Spin, swipe and tap to make robots move.
- Read up-to-date robotics news.
- Rate robots and check their ranking.
- View photography, videos and technical specs.
- Play *Faceoff*, an interactive question game.

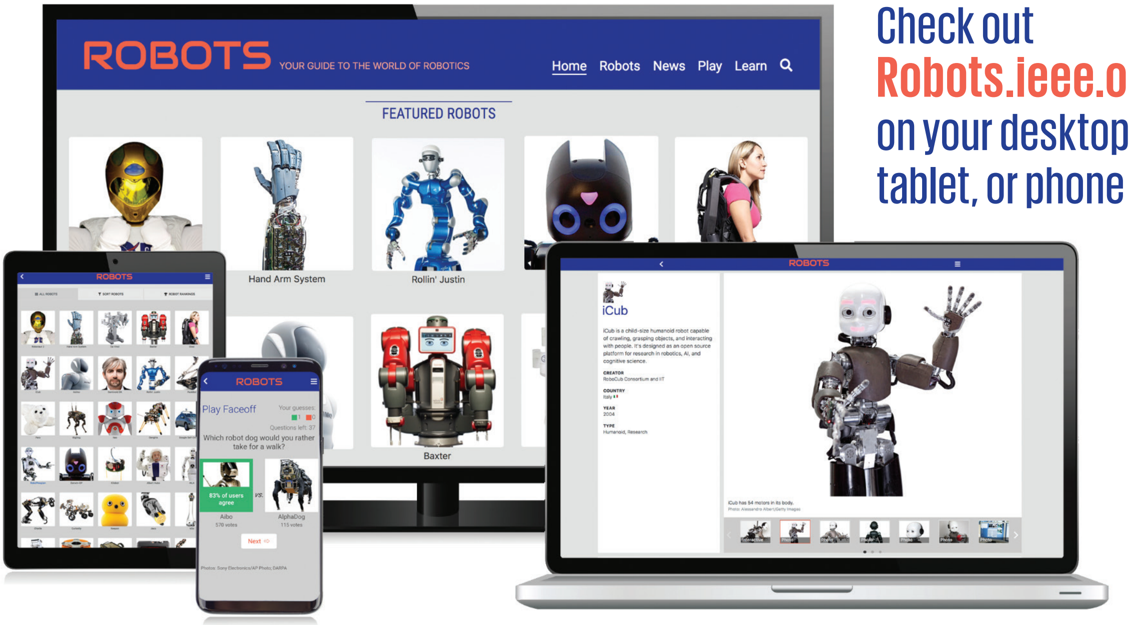

Play Faceoff

Your guesses: 🟩 1 🟥 0

Questions left: 37

Which robot dog would you rather take for a walk?

83% of users agree

Aibo
570 votes

vs.

AlphaDog
115 votes

Photos: Sony Electronics/AP Photo; DARPA

iCub

iCub is a child-size humanoid robot capable of crawling, grasping objects, and interacting with people. It's designed as an open source platform for research in robotics, AI, and cognitive science.

CREATOR
RoboCub Consortium and IIT

COUNTRY
Italy

YEAR
2004

TYPE
Humanoid, Research

Check out
Robots.ieee.org
on your desktop, tablet, or phone now!

# The Other Side of the Innovator's Dilemma

Beware the quiet disruptors

**In 1997, Harvard Business School** professor Clayton Christensen created a sensation among venture capitalists and entrepreneurs with his book *The Innovator's Dilemma*. The lesson that most people remember from it is that a well-run business can't afford to switch to a new approach—one that ultimately will replace its current business model—until it is too late.

One of the most famous examples of this conundrum involved photography. The large, very profitable companies that made film for cameras knew in the mid-1990s that digital photography would be the future, but there was never really a good time for them to make the switch. At almost any point they would have lost money. So what happened, of course, was that they were displaced by new companies making digital cameras. (Yes, Fujifilm did survive, but the transition was not pretty, and it involved an improbable series of events, machinations, and radical changes.)

A second lesson from Christensen's book is less well remembered but is an integral part of the story. The new companies springing up might get by for years with a disastrously less capable technology. Some of them, nevertheless, survive by finding a new niche they can fill that the incumbents cannot. That is where they quietly grow their capabilities.

For example, the early digital cameras had much lower resolution than film cameras, but they were also much smaller. I used to carry one on my key chain in my pocket and take photos of the participants in every meeting I had. The resolution was way too low to record stunning vacation vistas, but it was good enough to augment my poor memory for faces.

This lesson also applies to research. A great example of an underperforming new approach was the second wave of neural networks during the 1980s and 1990s that would eventually revolutionize artificial intelligence starting around 2010.

> **In 2012, the poor cousin of computer vision triumphed, and it completely changed the field of AI.**

Neural networks of various sorts had been studied as mechanisms for machine learning since the early 1950s, but they weren't very good at learning interesting things.

In 1979, Kunihiko Fukushima first published his research on something he called shift-invariant neural networks, which enabled his self-organizing networks to learn to classify handwritten digits wherever they were in an image. Then, in the 1980s, a technique called backpropagation was rediscovered; it allowed for a form of supervised learning in which the network was told what the right answer should be. In 1989, Yann LeCun combined backpropagation with Fuksuhima's ideas into something that has come to be known as convolutional neural networks (CNNs). LeCun, too, concentrated on images of handwritten digits.

Over the next 10 years, the U.S. National Institute of Standards and Technology (NIST) came up with a database, which was modified by LeCun, consisting of 60,000 training digits and 10,000 test digits. This standard test database, called MNIST, allowed researchers to precisely measure and compare the effectiveness of different improvements to CNNs. There was a lot of progress, but CNNs were no match for the entrenched AI methods in computer vision when applied to arbitrary images generated by early self-driving cars or industrial robots.

But during the 2000s, more and more learning techniques and algorithmic improvements were added to CNNs, leading to what is now known as deep learning. In 2012, suddenly, and seemingly out of nowhere, deep learning outperformed the standard computer-vision algorithms in a set of test images of objects, known as ImageNet. The poor cousin of computer vision triumphed, and it completely changed the field of AI.

A small number of people had labored for decades and surprised everyone. Congratulations to

all of them, both well known and not so well known.

But beware. The message of Christensen's book is that such disruptions never stop. Those standing tall today will be surprised by new methods that they have not begun to consider. There are small groups of renegades trying all sorts of new things, and some of them, too, are willing to labor quietly and against all odds for decades. One of those groups will someday surprise us all.

I love this aspect of technological and scientific disruption. It is what makes us humans great. And dangerous. ■

*Photo-illustration by Edmon de Haro*

AI-based systems guide robotic arms to grab materials from a stream of mixed recyclables and place them in the correct bins. Here, a tandem robot system operates at a Waste Connections recycling facility [below], and a single robot arm [right] recovers a piece of corrugated cardboard. The United States does a pretty good job when it comes to cardboard: In 2021, 91.4 percent of discarded cardboard was recycled, according to the American Forest and Paper Association.

ALL PHOTOS: AMP ROBOTICS

*Computer-vision systems sort your recyclables at superhuman speed*

***By Jason Calaiaro***

# AI Takes a Dumpster Dive

**It's Tuesday night.** In front of your house sits a large blue bin, full of newspaper, cardboard, bottles, cans, foil take-out trays, and empty yogurt containers. You may feel virtuous, thinking you're doing your part to reduce waste. But after you rinse out that yogurt container and toss it into the bin, you probably don't think much about it ever again.

The truth about recycling in many parts of the United States and much of Europe is sobering. Tomorrow morning, the contents of the recycling bin will be dumped into a truck and taken to the recycling facility to be sorted. Most of the material will head off for processing and eventual use in new products. But a lot of it will end up in a landfill.

So how much of the material that goes into the typical bin avoids a trip to landfill? For countries that do curbside recycling, the number—called the recovery rate—appears to average around 70 to 90 percent, though widespread data isn't available. That doesn't seem bad. But in some municipalities, it can go as low as 40 percent.

What's worse, only a small quantity of all recyclables makes it into the bins—just 32 percent in the United States and 10 to 15 percent globally. That's a lot of material made from finite resources that needlessly goes to waste.

We have to do better than that. Right now, the recycling industry is facing a financial crisis, thanks to falling prices for sorted recyclables as well as policy,

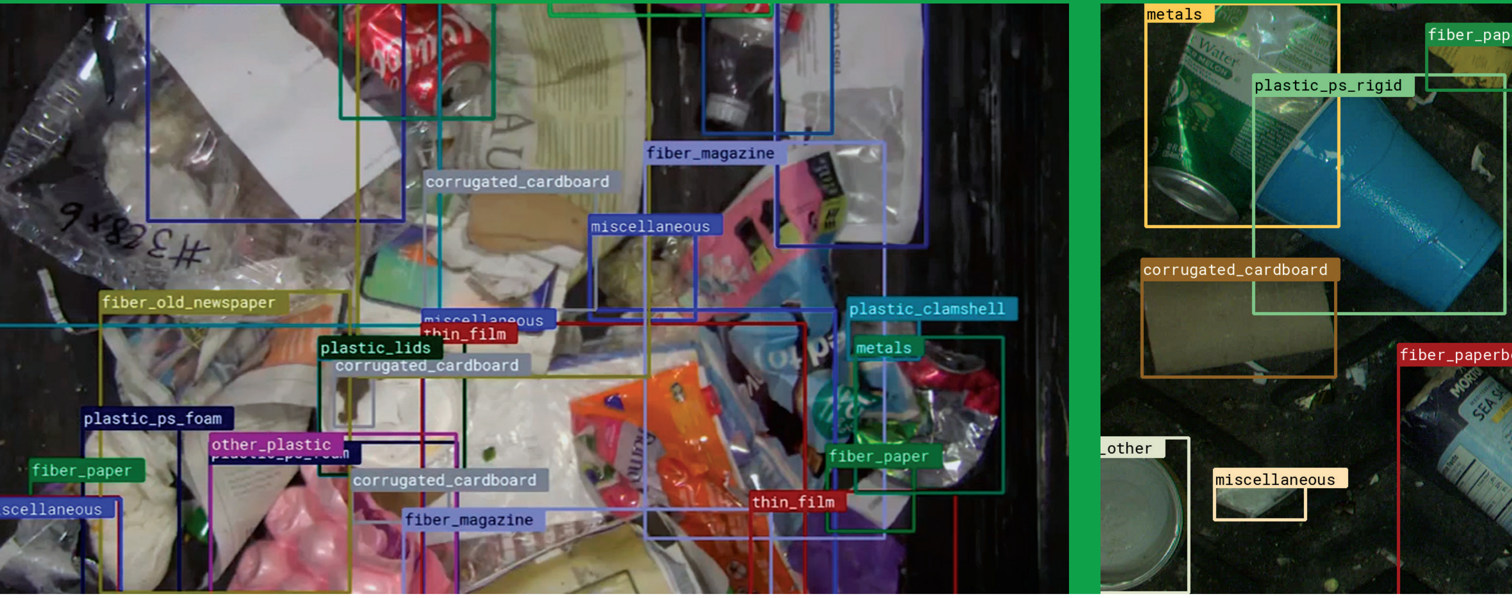

enacted by China in 2018, which restricts the import of many materials destined for recycling and shuts out most recyclables originating in the United States.

There is a way to do better. Using computer vision, machine learning, and robots to identify and sort recycled material, we can improve the accuracy of automatic sorting machines, reduce the need for human intervention, and boost overall recovery rates.

My company, Amp Robotics, based in Louisville, Colo., is developing hardware and software that relies on image analysis to sort recyclables with far higher accuracy and recovery rates than are typical for conventional systems. Other companies are similarly working to apply AI and robotics to recycling, including Bulk Handling Systems, Machinex, and Tomra. To date, the technology has been installed in hundreds of sorting facilities around the world. Expanding its use will prevent waste and help the environment by keeping recyclables out of landfills and making them easier to reprocess and reuse.

**Before I explain how AI will improve recycling,** let's look at how recycled materials were sorted in the past and how they're being sorted in most parts of the world today.

When recycling began in the 1960s, the task of sorting fell to the consumer—newspapers in one bundle, cardboard in another, and glass and cans in their own separate bins. That turned out to be too much of a hassle for many people and limited the amount of recyclable materials gathered.

In the 1970s, many cities took away the multiple bins and replaced them with a single container, with sorting happening downstream. This "single stream" recycling boosted participation, and it is now the dominant form of recycling in developed countries.

Moving the task of sorting further downstream led to the building of sorting facilities. To do the actual sorting, recycling entrepreneurs adapted equipment from the mining and agriculture industries, filling in with human labor as necessary. These sorting systems had no computer intelligence, relying instead on the physical properties of materials to separate them. Glass, for example, can be broken into tiny pieces and then sifted and collected. Cardboard is rigid and light—it can glide over a series of mechanical camlike disks, while other, denser materials fall in between the disks. Ferrous metals can be magnetically separated from other materials; magnetism can also be induced in nonferrous items, like aluminum, using a large eddy current.

By the 1990s, hyperspectral imaging, developed by NASA and first launched in a satellite in 1972, was becoming commercially viable and began to show up in the recycling world. Unlike human eyes, which mostly see in combinations of red, green, and blue, hyperspectral sensors divide images into many more spectral bands. The technology's ability to distinguish between different types of plastics changed the game for recyclers, bringing not only optical sensing but computer intelligence into the process. Programmable optical sorters were also developed to separate paper products, distinguishing, say, newspaper from junk mail.

So today, much of the sorting is automated. These systems generally sort to 80 to 95 percent purity—that is, 5 to 20 percent of the output shouldn't be there. For the output to be profitable, however, the purity must be higher than 95 percent; below this threshold, the value drops, and often it's worth nothing. So humans manually clean up each of the streams, picking out stray objects before the material is compressed and baled for shipping.

Despite all the automated and manual sorting, about 10 to 30 percent of the material that enters the facility ultimately ends up in a landfill. In most cases, more than half of that material is recyclable and worth money but was simply missed.

We've pushed the current systems as far as they can go. Only AI can do better.

**Getting AI into the recycling business** means combining pick-and-place robots with accurate real-time object detection. Pick-

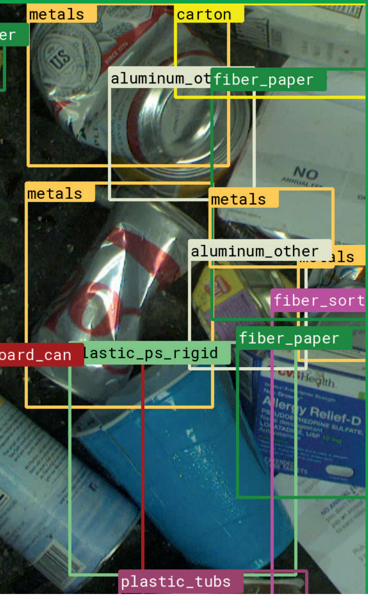

The Amp Cortex, a high-speed robotic sorting system guided by artificial intelligence, identifies materials by category on a conveyor belt. To date, systems in operation have recognized more than 50 billion objects in various permutations.

and-place robots combined with computer vision systems are used in manufacturing to grab particular objects, but they generally are just looking repeatedly for a single item, or for a few items of known shapes and under controlled lighting conditions. Recycling, though, involves infinite variability in the kinds, shapes, and orientations of the objects traveling down the conveyor belt, requiring nearly instantaneous identification along with the quick dispatch of a new trajectory to the robot arm.

My company first began using AI in 2016 to extract empty cartons from other recyclables at a facility in Colorado; today, we have systems installed in more than 25 U.S. states and six countries. We weren't the first company to try AI sorting, but it hadn't previously been used commercially. And we have steadily expanded the types of recyclables our systems can recognize and sort.

AI makes it theoretically possible to recover all of the recyclables from a mixed-material stream at accuracy approaching 100 percent, entirely based on image analysis. If an AI-based sorting system can see an object, it can accurately sort it.

Consider a particularly challenging material for today's recycling sorters: high-density polyethylene (HDPE), a plastic commonly used for detergent bottles and milk jugs. (In the United States, Europe, and China, HDPE products are labeled as No. 2 recyclables.) In a system that relies on hyperspectral imaging, batches of HDPE tend to be mixed with other plastics and may have paper or plastic labels, making it difficult for the hyperspectral imagers to detect the underlying object's chemical composition.

An AI-driven computer-vision system, by contrast, can determine that a bottle is HDPE and not something else by recognizing its packaging. Such a system can also use attributes like color, opacity, and form factor to increase detection accuracy, and even sort by color or specific product, reducing the amount of reprocessing needed. Though the system doesn't attempt to understand the meaning of words on labels, the words are part of an item's visual attributes.

We at Amp Robotics have built systems that can do this kind of sorting. In the future, AI systems could also sort by combinations of material and by original use, enabling food-grade materials to be separated from containers that held household cleaners, and paper contaminated with food waste to be separated from clean paper.

**Training a neural network** to detect objects in the recycling stream is not easy. It is at least several orders of magnitude more challenging than recognizing faces in a photograph, because there can be a nearly infinite variety of ways that recyclable materials can be deformed, and the system has to recognize the permutations.

It's hard enough to train a neural network to identify all the different types of bottles of laundry detergent on the market today, but it's an entirely different challenge when you consider the physical deformations that these objects can undergo by the time they reach a recycling facility. They can be folded, torn, or smashed. Mixed into a stream of other objects, a bottle might have only a corner visible. Fluids or food waste might obscure the material.

We train our systems by giving them images of materials belonging to each category, sourced from recycling facilities around the world. My company now has the world's largest data set of recyclable material images for use in machine learning.

Using this data, our models learn to identify recyclables in the same way their human counterparts do, by spotting patterns and features that distinguish different materials. We continu-

An AI-guided robotic system, one of several installed at a single facility, identifies and recovers mixed plastics. The difficulty of sorting and recycling plastics drives down overall recycling rates. But AI-enabled sorting that can separate plastics by color, clarity, and other features could dramatically change the amount of plastics that get reprocessed instead of landfilled or combusted.

# INSIDE THE SORTING CENTER

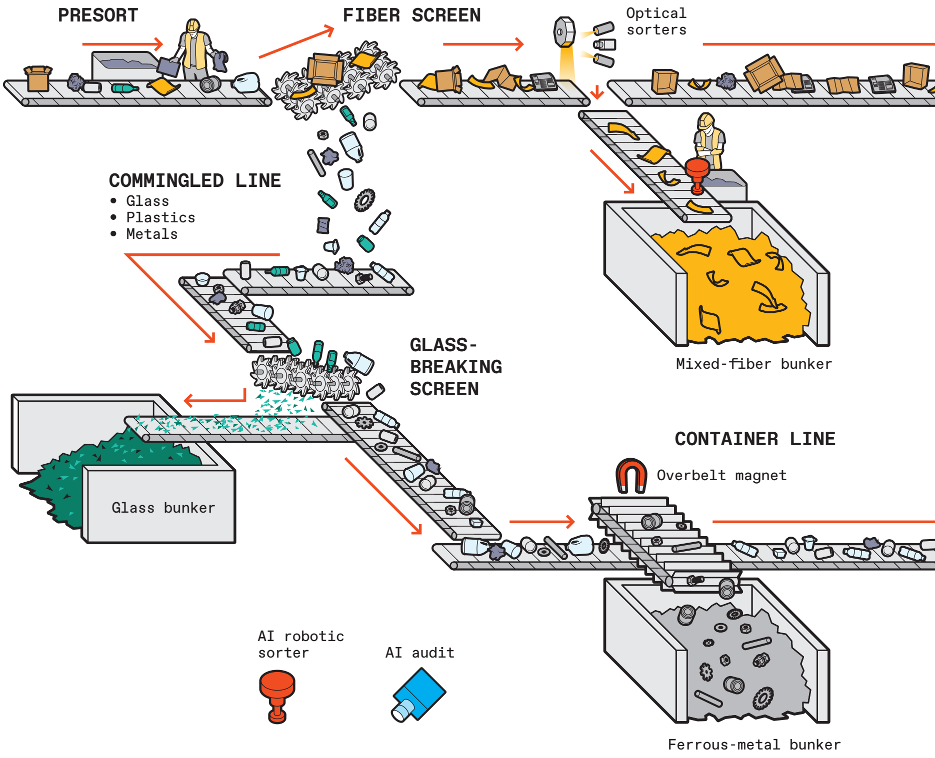

Today's recycling facilities use mechanical sorting, optical hyperspectral sorting, and human workers. Here's what typically happens after the recycling truck leaves your house with the contents of your blue bin.

Trucks unload on a concrete pad, called the tip floor. A front-end loader scoops up material in bulk and dumps it onto a conveyor belt, typically at a rate of 30 to 60 tonnes per hour.

The first stage is the presort. Human workers remove large or problematic items that shouldn't have made it onto collection trucks in the first place—bicycles, big pieces of plastic film, propane canisters, car transmissions.

Sorting machines that rely on optical hyperspectral imaging or human workers separate fiber (office paper, cardboard, magazines—referred to as 2D products, as they are mostly flat) from the remaining plastics and metals. In the case of the optical sorters, cameras stare down at the material rolling down the conveyor belt, detect an object made of the target substance, and then send a message to activate a bank of electronically controllable solenoids to divert the object into a collection bin.

The nonfiber materials pass through a mechanical system with densely packed camlike wheels. Large items glide past while small items, like that recyclable fork you thoughtfully deposited in your blue bin, slip through, headed straight for landfill—they are just too small to be sorted. Machines also smash glass, which falls to the bottom and is screened out.

The rest of the stream then passes under overhead magnets, which collect items made of ferrous metals, and an eddy-current-inducing machine, which jolts nonferrous metals to another collection area.

At this point, mostly plastics remain. More hyperspectral sorters, in series, can pull off plastics one type—like the HDPE of detergent bottles or the PET of water bottles—at a time.

ously collect random samples from all the facilities that use our systems, and then annotate them, add them to our database, and retrain our neural networks. We also test our networks to find models that perform best on target material and do targeted additional training on materials that our systems have trouble identifying correctly.

In general, neural networks are susceptible to learning the wrong thing. Pictures of cows are associated with milk packaging, which is commonly produced as a fiber carton or HDPE container. But milk products can also be packaged in other plastics; for example, single-serving milk bottles may look like the HDPE of gallon jugs but are usually made from an opaque form of the PET (polyethylene terephthalate) used for water bottles. Cows don't always mean fiber or HDPE, in other words.

There is also the challenge of staying up to date with the continual changes in consumer packaging. Any mechanism that relies on visual observation to learn associations between packaging and material types will need to consume a steady stream of data to ensure that objects are classified accurately.

But we can get these systems to work. Right now, our systems do really well on certain categories—more than 98 percent accuracy on aluminum cans—and are getting better at distinguishing nuances like color, opacity, and initial use (spotting those food-grade plastics).

**Now that AI-based systems** are ready to take on your recyclables, how might things change? Certainly, they will boost the use of robotics, which is only minimally used in the recycling industry today. Given the perpetual worker shortage in this dull and dirty business, automation is a path worth taking.

AI can also help us understand how well today's existing sorting processes are doing and how we can improve them. Today, we have a very crude understanding of the operational efficiency of sorting facilities—we weigh trucks on the way in and weigh the output on the way out. No facility can tell you the purity of the products with any certainty; they only audit quality periodically by breaking open random bales. But if you placed an AI-powered vision system over the inputs and outputs of relevant parts of the sorting process, you'd gain a holistic view of what material is flowing where. This level of scrutiny is just beginning in hundreds of facilities around the world, and it should lead to greater efficiency in recycling operations. Being able to digitize the real-time flow

*Illustration by Chris Philpot*

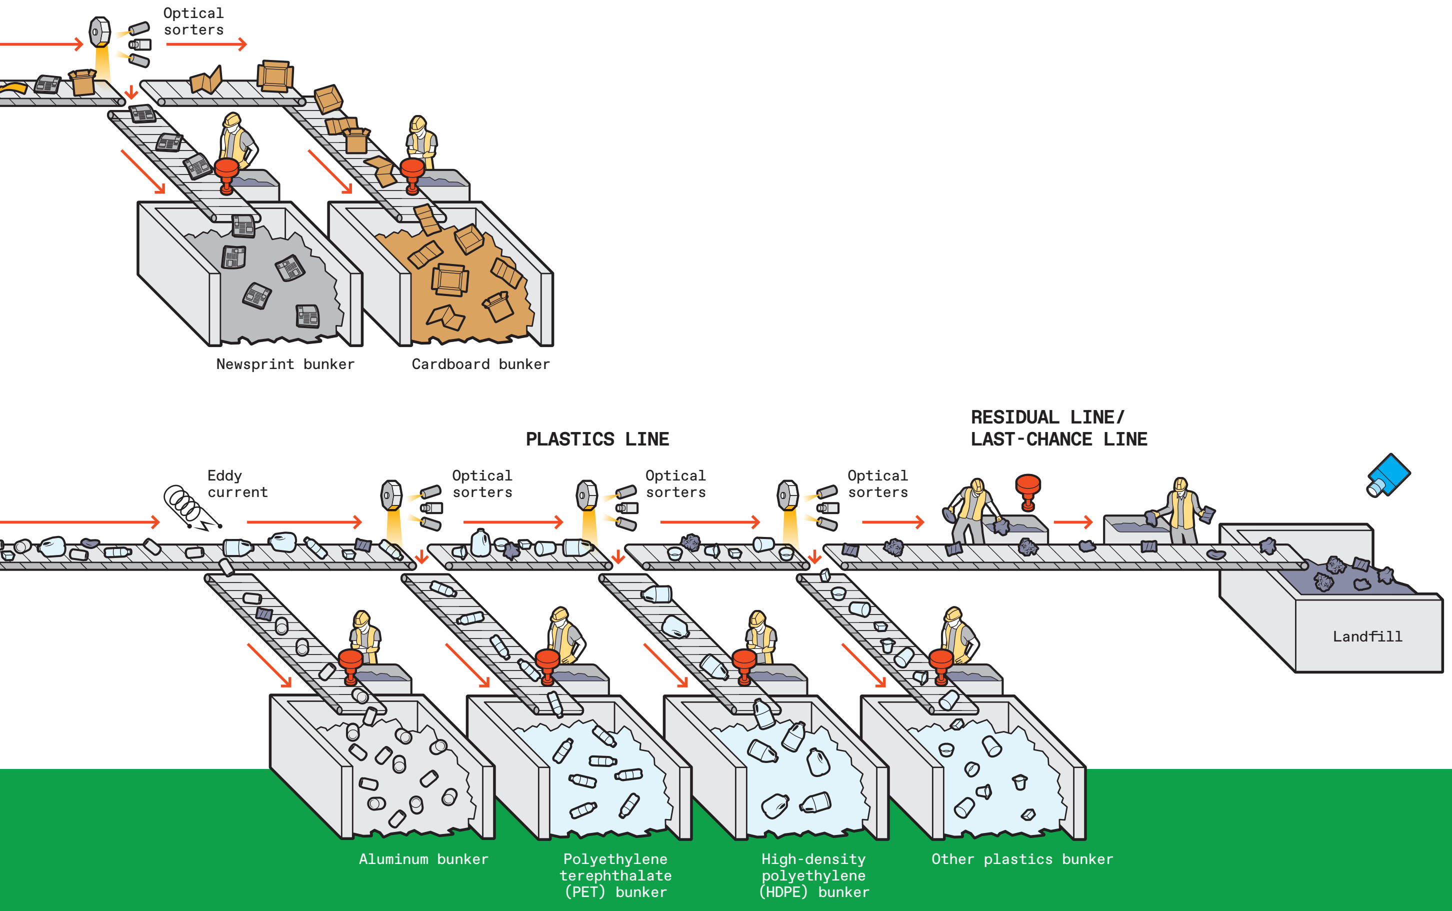

Finally, whatever is left—between 10 to 30 percent of what came in on the trucks—goes to landfill.

In the future, AI-driven robotic sorting systems and AI inspection systems could replace human workers at most points in this process. In the diagram, red icons indicate where AI-driven robotic systems could replace human workers and a blue icon indicates where an AI auditing system could make a final check on the success of the sorting effort.

of recyclables with precision and consistency also provides opportunities to better understand which recyclable materials are and are not currently being recycled and then to identify gaps that will allow facilities to improve their recycling systems overall.

But to really unleash the power of AI on the recycling process, we need to rethink the entire sorting process. Today, recycling operations typically whittle down the mixed stream of materials to the target material by removing nontarget material—they do a "negative sort," in other words. Instead, using AI vision systems with robotic pickers, we can perform a "positive sort." Instead of removing nontarget material, we identify each object in a stream and select the target material.

To be sure, our recovery rate and purity are only as good as our algorithms. Those numbers continue to improve as our systems gain more experience in the world and our training data set continues to grow. We expect to eventually hit purity and recovery rates of 100 percent.

The implications of moving from more mechanical systems to AI are profound. Rather than coarsely sorting to 80 percent purity and then manually cleaning up the stream to 95 percent purity, a facility can reach the target purity on the first pass. And instead of having a unique sorting mechanism handling each type of material, a sorting machine can change targets just by a switch in algorithm.

The use of AI also means that we can recover materials long ignored for economic reasons. Until now, it was only economically viable for facilities to pursue the most abundant, high-value items in the waste stream. But with machine-learning systems that do positive sorting on a wider variety of materials, we can start to capture a greater diversity of material at little or no overhead to the business. That's good for the planet.

We are beginning to see a few AI-based secondary recycling facilities go into operation, with Amp's technology first coming online in Denver in late 2020. These systems are currently used where material has already passed through a traditional sort, seeking high-value materials missed or low-value materials that can be sorted in novel ways and therefore find new markets.

Thanks to AI, the industry is beginning to chip away at the mountain of recyclables that end up in landfills each year—a mountain containing billions of tons of recyclables representing billions of dollars lost and nonrenewable resources wasted. ■

# Quantum Error Correction at the Threshold

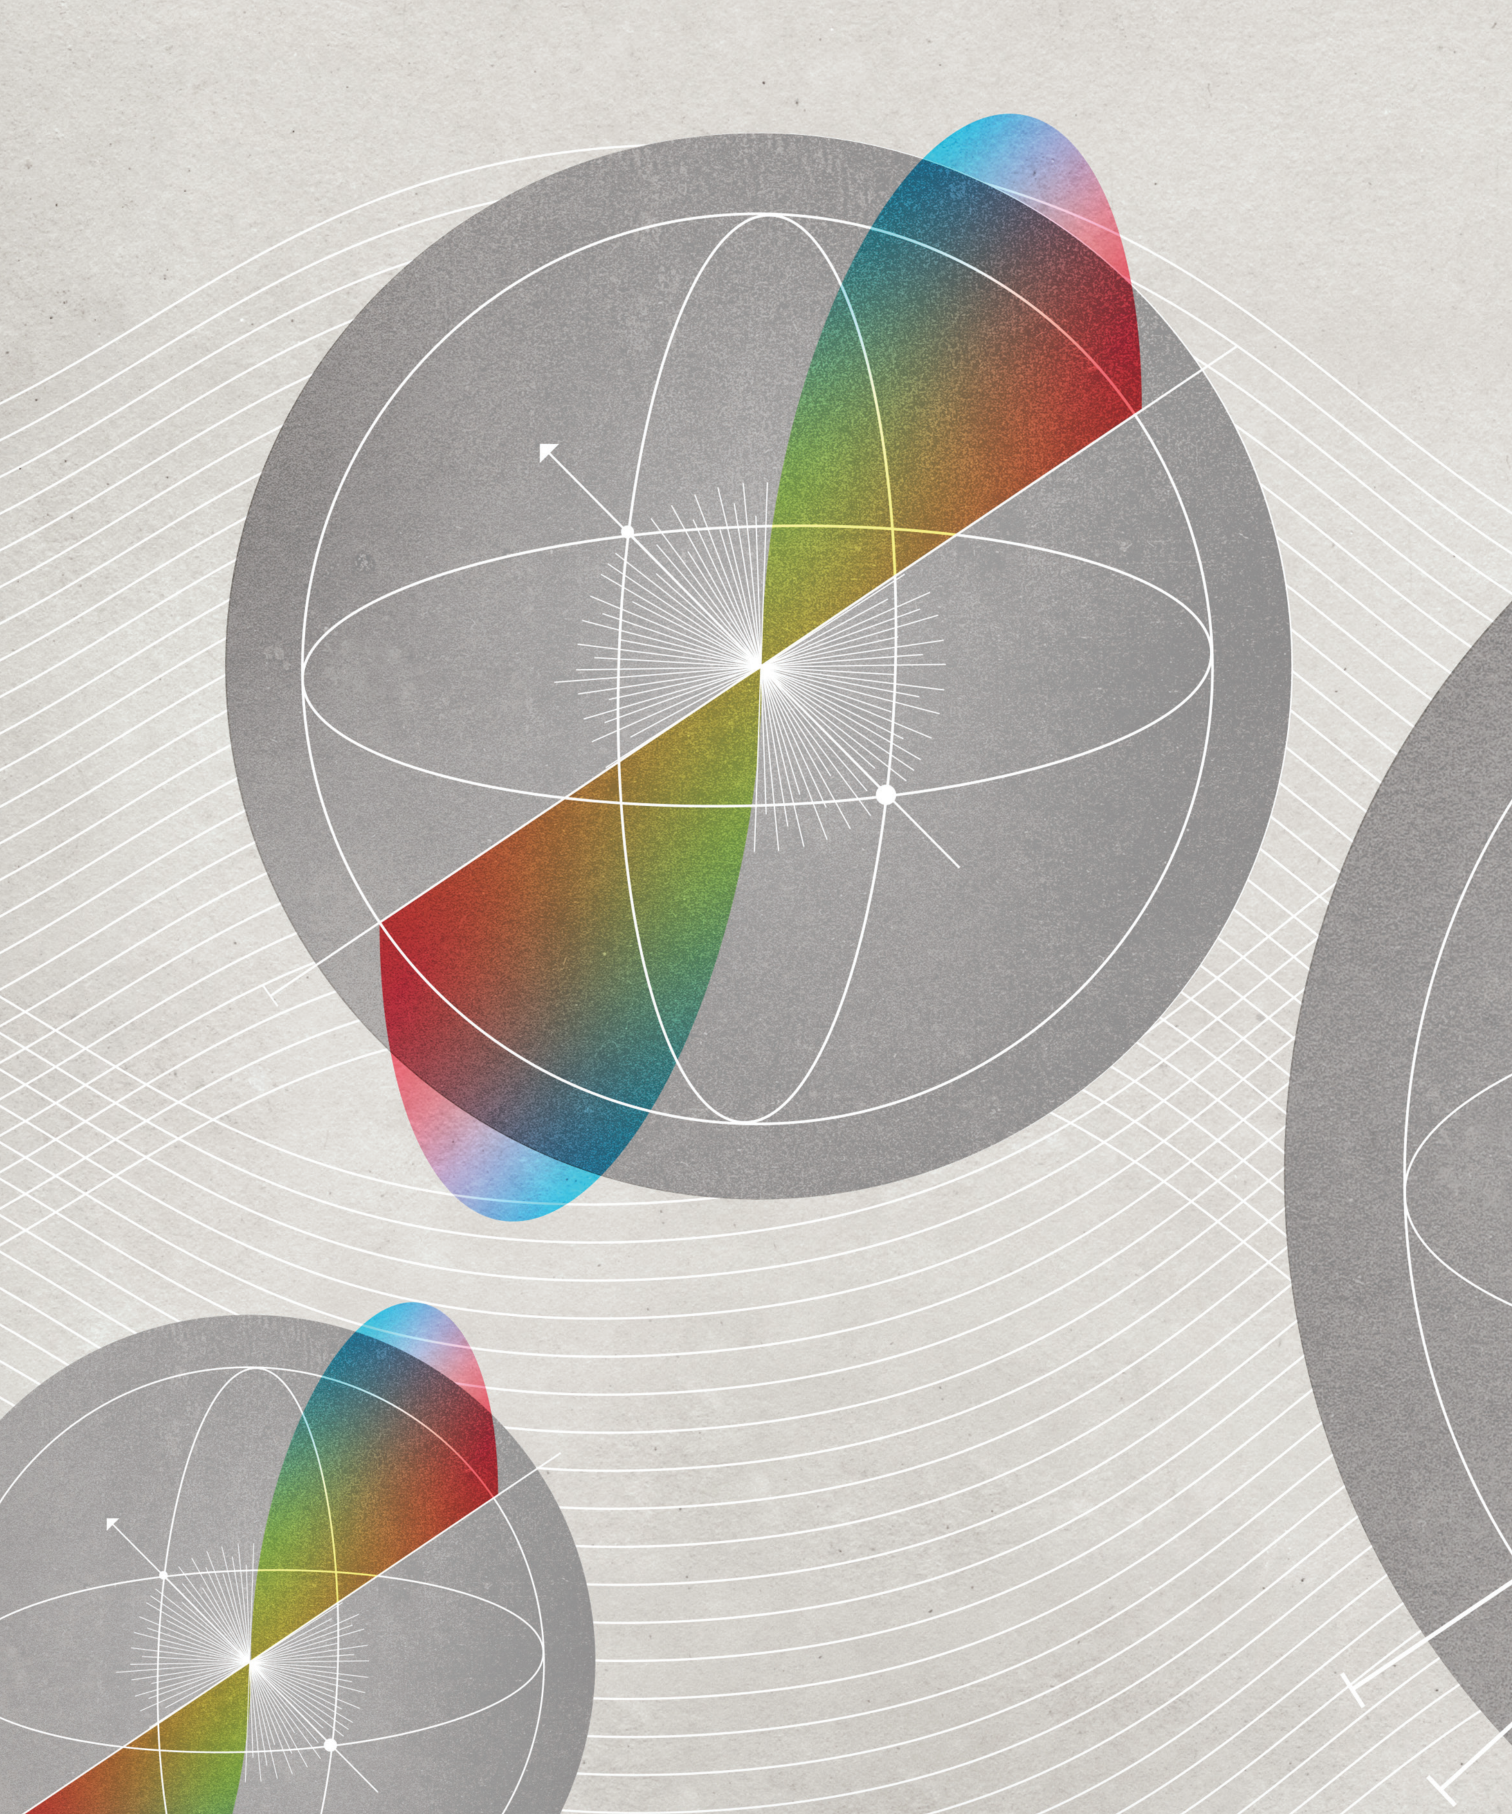

If technologists don't get beyond it, quantum computers will never be big

BY MICHAEL J. BIERCUK & THOMAS M. STACE

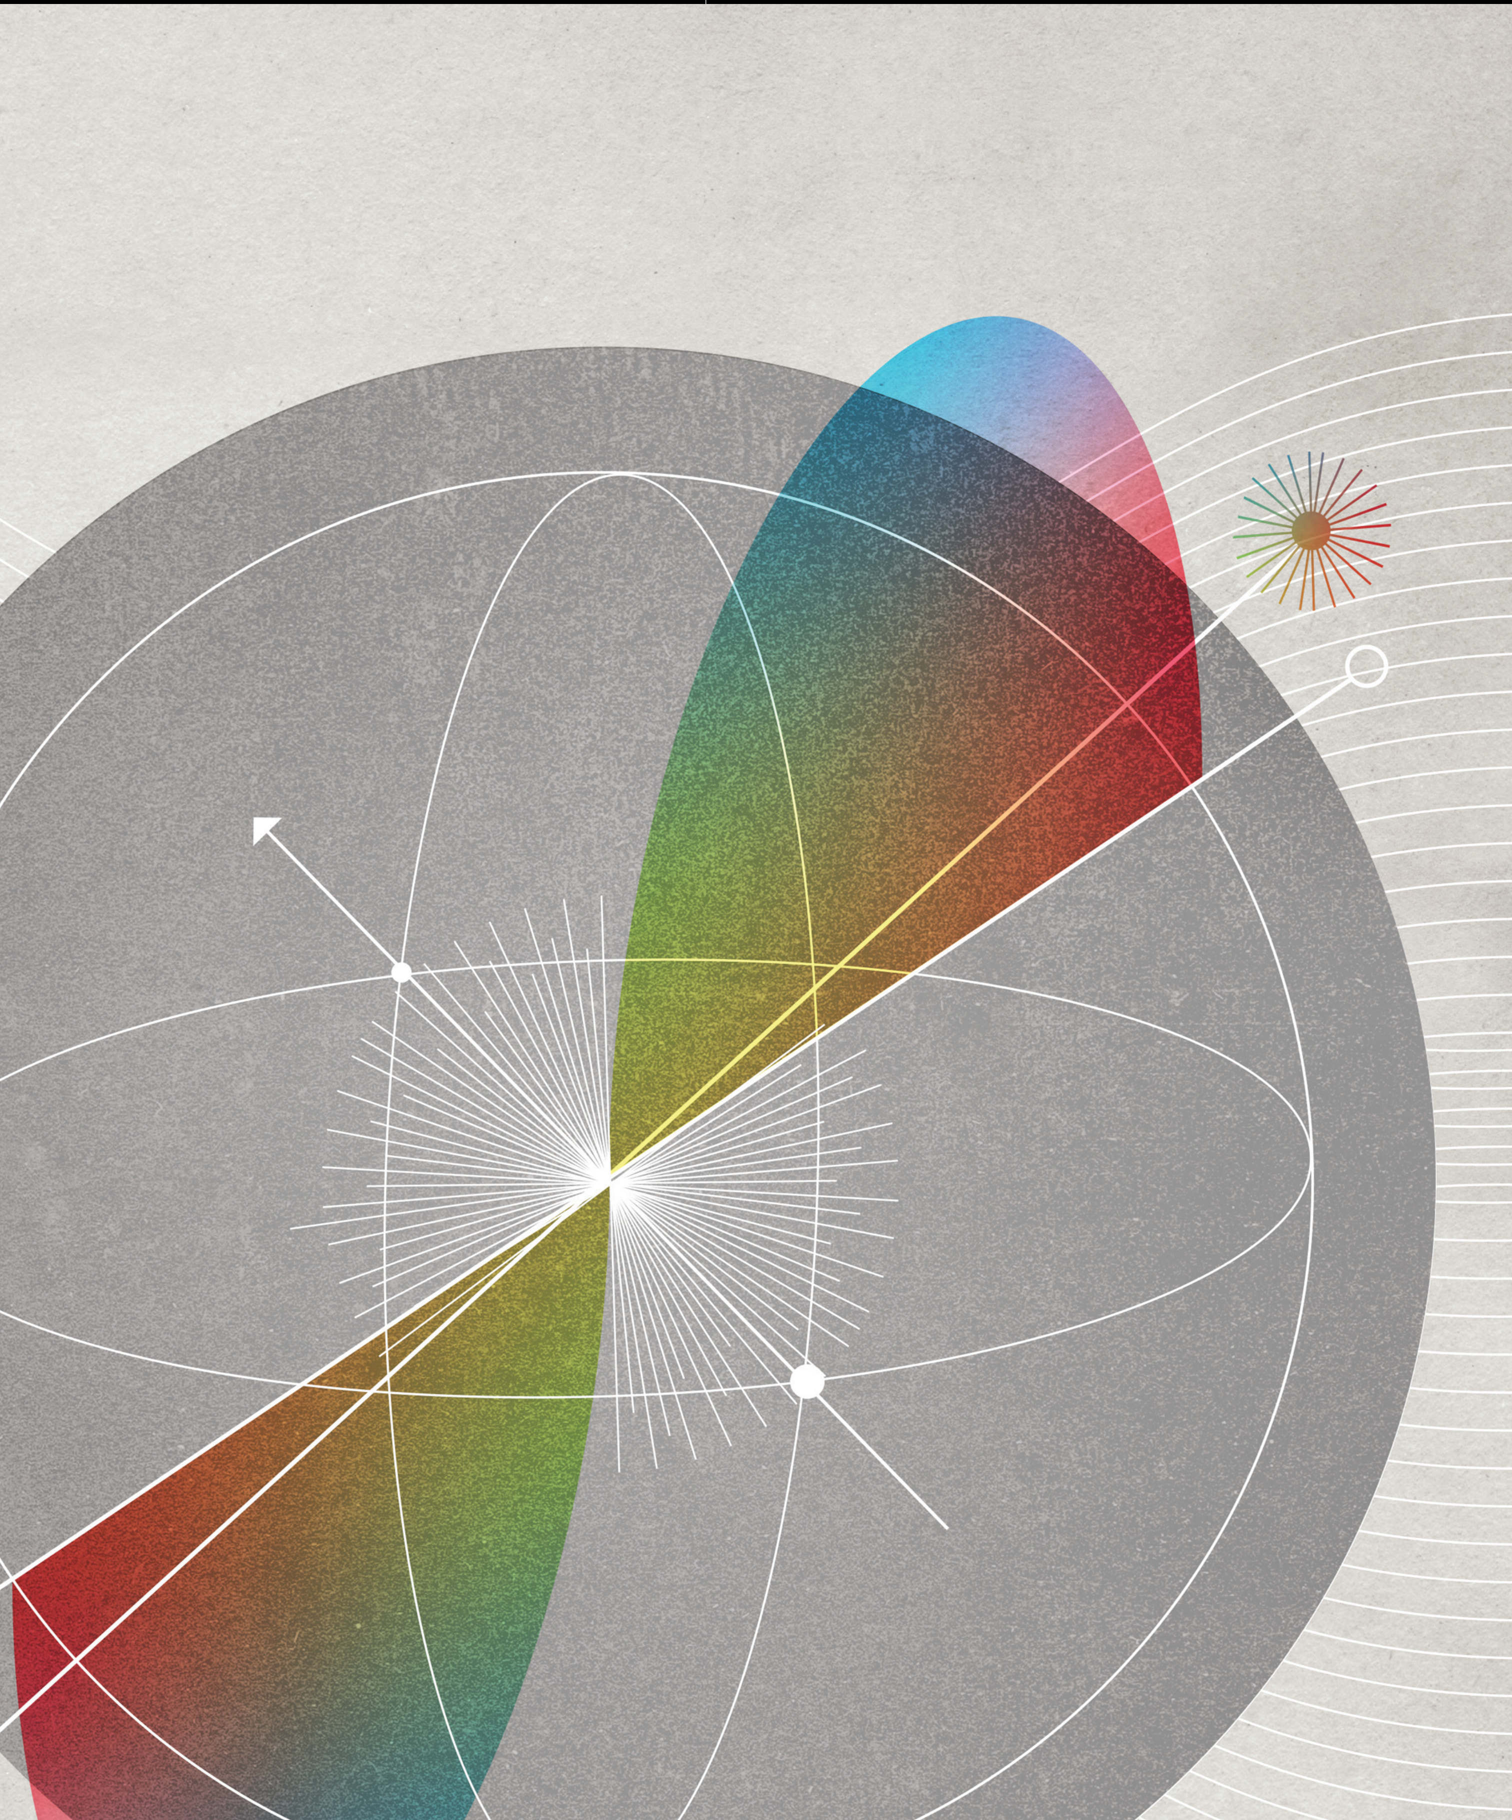

**DATES CHISELED** into an ancient tombstone have more in common with the data in your phone or laptop than you may realize. They both involve conventional, classical information, carried by hardware that is relatively immune to errors. The situation inside a quantum computer is far different: The information itself has its own idiosyncratic properties, and compared with standard digital microelectronics, state-of-the-art quantum-computer hardware is more than a billion trillion times as likely to suffer a fault. This tremendous susceptibility to errors is the single biggest problem holding back quantum computing from realizing its great promise.

Fortunately, an approach known as quantum error correction (QEC) can remedy this problem, at least in principle. A mature body of theory built up over the past quarter century now provides a solid theoretical foundation, and experimentalists have demonstrated dozens of proof-of-principle examples of QEC. But these experiments still have not reached the level of quality and sophistication needed to reduce the overall error rate in a system.

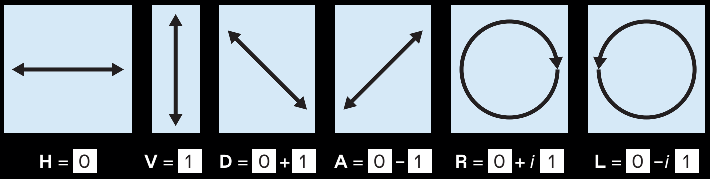

Polarized light is an example of superposition. A classical binary digit could be represented by encoding 0 as horizontally (H) polarized light, and 1 as vertically (V) polarized light. Light polarized at other angles has components of both H and V, representing 0 and 1 simultaneously. Examples include the diagonal (D) polarization at 45°, the antidiagonal (A) at −45°, as well as right (R) and left (L) circularly polarized light (the imaginary number *i* represents a difference in phase). These states become fully fledged quantum bits (qubits) when they consist of pulses that each contain a single photon.

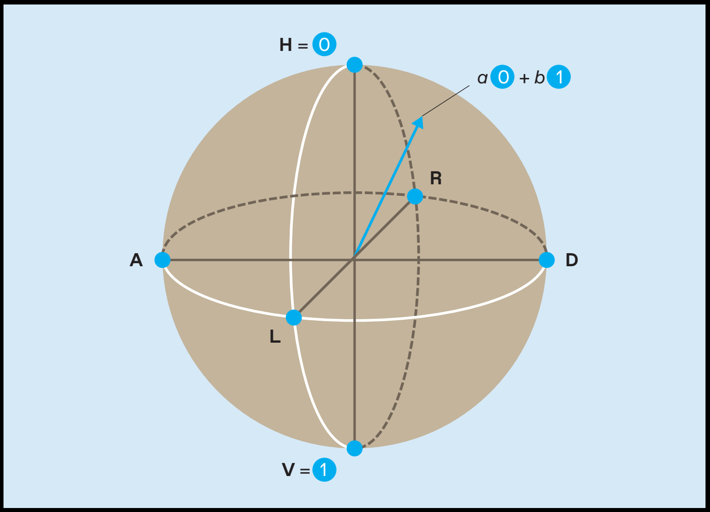

The possible states of a single isolated qubit [blue arrow] are neatly represented on a sphere, known as a Bloch sphere. The states 0 and 1 sit at the north and south poles, and the polarization states D, A, R, and L lie on the equator. Other possible superpositions of 0 and 1 (described by complex numbers *a* and *b*) cover the rest of the surface. Noise can make the qubit state wander continuously from its correct location.

The two of us, along with many other researchers involved in quantum computing, are trying to move definitively beyond these preliminary demos of QEC so that it can be employed to build useful, large-scale quantum computers. But before describing how we think such error correction can be made practical, we need to first review what makes a quantum computer tick.

**INFORMATION IS PHYSICAL.** This was the mantra of the distinguished IBM researcher Rolf Landauer. Abstract though it may seem, information always involves a physical representation, and the physics matters.

Conventional digital information consists of bits, zeros and ones, which can be represented by classical states of matter, that is, states well described by classical physics. Quantum information, by contrast, involves *qubits*—quantum bits—whose properties follow the peculiar rules of quantum mechanics.

A classical bit has only two possible values: 0 or 1. A qubit, however, can occupy a superposition of these two information states, taking on characteristics of both. Polarized light provides intuitive examples of superpositions. You could use horizontally polarized light to represent 0 and vertically polarized light to represent 1, but light can also be polarized on an angle and then has both horizontal and vertical components at once. Indeed, one way to represent a qubit is by the polarization of a single photon of light.

These ideas generalize to groups of $n$ bits or qubits: $n$ bits can represent any one of $2^n$ possible values at any moment, while $n$ qubits can include components corresponding to all $2^n$ classical states simultaneously in superposition. These superpositions provide a vast range of possible states for a quantum computer to work with, albeit with limitations on how they can be manipulated and accessed. Superposition of information is a central resource used in quantum processing and, along with other quantum rules, enables powerful new ways to compute.

Researchers are experimenting with many different physical systems to hold and process quantum information, including light, trapped atoms and ions, and solid-state devices based on semiconductors or superconductors. For the purpose of realizing qubits, all these systems follow the same underlying mathe-

*Previous pages: Illustration by Chad Hagen*

matical rules of quantum physics, and all of them are highly sensitive to environmental fluctuations that introduce errors. By contrast, the transistors that handle classical information in modern digital electronics can reliably perform a billion operations per second for decades with a vanishingly small chance of a hardware fault.

Of particular concern is the fact that qubit states can roam over a continuous range of superpositions. Polarized light again provides a good analogy: The angle of linear polarization can take *any* value from 0 to 180 degrees.

Pictorially, a qubit's state can be thought of as an arrow pointing to a location on the surface of a sphere. Known as a Bloch sphere, its north and south poles represent the binary states 0 and 1, respectively, and all other locations on its surface represent possible quantum superpositions of those two states. Noise causes the Bloch arrow to drift around the sphere over time. A conventional computer represents 0 and 1 with physical quantities, such as capacitor voltages, that can be locked near the correct values to suppress this kind of continuous wandering and unwanted bit flips. There is no comparable way to lock the qubit's "arrow" to its correct location on the Bloch sphere.

Early in the 1990s, Landauer and others argued that this difficulty presented a fundamental obstacle to building useful quantum computers. The issue is known as scalability: Although a simple quantum processor performing a few operations on a handful of qubits might be possible, could you scale up the technology to systems that could run lengthy computations on large arrays of qubits? A type of classical computation called analog computing also uses continuous quantities and is suitable for some tasks, but the problem of continuous errors prevents the complexity of such systems from being scaled up. Continuous errors with qubits seemed to doom quantum computers to the same fate.

We now know better. Theoreticians have successfully adapted the theory of error correction for classical digital data to quantum settings. QEC makes scalable quantum processing possible in a way that is impossible for analog computers. To get a sense of how it works, it's worthwhile to review how error correction is performed in classical settings.

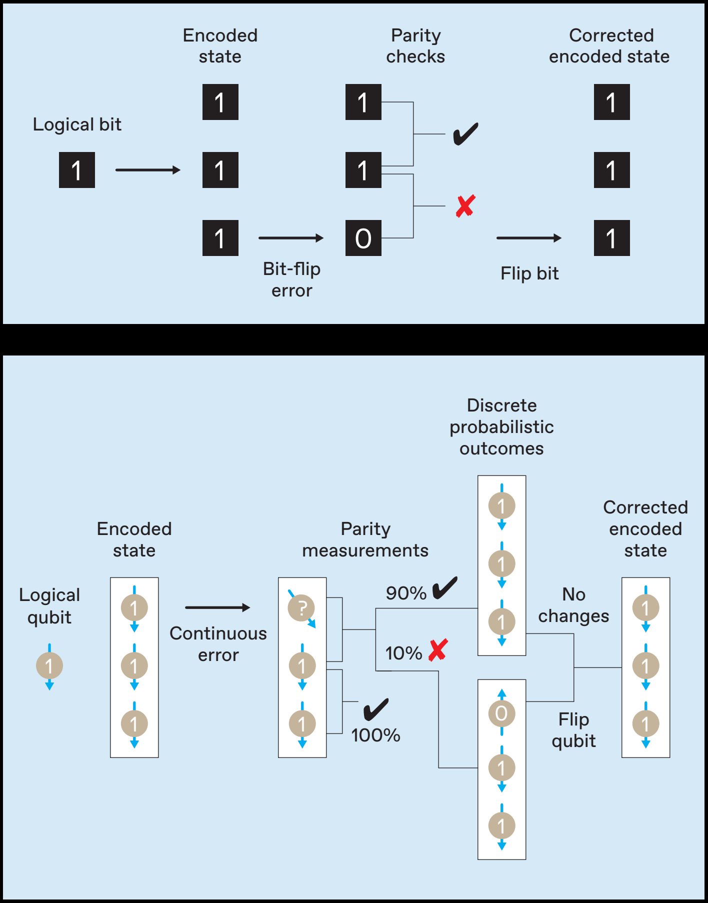

**Simple repetition code [top] on a conventional bit allows single bit-flip errors to be detected via parity checks and then corrected. A similar code for qubits [bottom] must deal with continuous errors. (For simplicity, we depict the case of a logical qubit in a nonsuperposition state, 1.) The parity checks, being quantum measurements, produce discrete outcomes with various probabilities, converting the continuous error into a discrete one and allowing correction by a qubit flip. The individual qubit states are not revealed by the parity measurements.**

**SIMPLE SCHEMES** can deal with errors in classical information. For instance, in the 19th century, ships routinely carried clocks for determining the ship's longitude during voyages. A good clock that could keep track of the time in Greenwich, in combination with the sun's position in the sky, provided the necessary data. A mistimed clock could lead to dangerous navigational errors, though, so ships often carried at least three of them. Two clocks reading different times could detect when one was at fault, but three were needed to identify which timepiece was faulty and correct it through a majority vote.

The use of multiple clocks is an example of a repetition code: Information is redundantly encoded in multiple physical devices such that a disturbance in one can be identified and corrected.

As you might expect, quantum mechanics adds some major complications when dealing with errors. Two problems in particular might seem to dash any hopes of using a quantum repetition code. The first problem is that measurements fundamentally disturb quantum systems. So if you encoded information on three

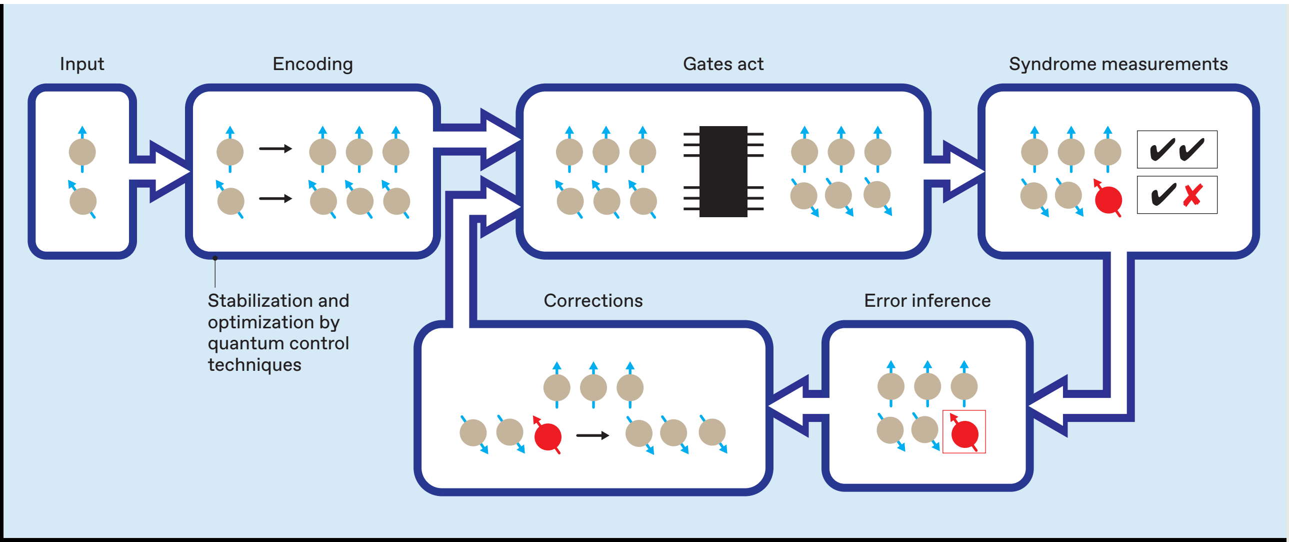

A long quantum computation will require many cycles of quantum error correction (QEC). Each cycle would consist of gates acting on encoded qubits (performing the computation), followed by syndrome measurements from which errors can be inferred, and corrections. The effectiveness of this QEC feedback loop can be greatly enhanced by including quantum-control techniques (represented by the thick blue outline) to stabilize and optimize each of these processes.

qubits, for instance, observing them directly to check for errors would ruin them. Like Schrödinger's cat when its box is opened, their quantum states would be irrevocably changed, spoiling the very quantum features your computer was intended to exploit.

The second issue is a fundamental result in quantum mechanics called the no-cloning theorem, which tells us it is impossible to make a perfect copy of an unknown quantum state. If you know the exact superposition state of your qubit, there is no problem producing any number of other qubits in the same state. But once a computation is running and you no longer know what state a qubit has evolved to, you cannot manufacture faithful copies of that qubit except by duplicating the entire process up to that point.

Fortunately, you can sidestep both of these obstacles. We'll first describe how to evade the measurement problem using the example of a classical three-bit repetition code. You don't actually need to know the state of every individual code bit to identify which one, if any, has flipped. Instead, you ask two questions: "Are bits 1 and 2 the same?" and "Are bits 2 and 3 the same?" These are called parity-check questions because two identical bits are said to have even parity, and two unequal bits have odd parity.

The two answers to those questions identify which single bit has flipped, and you can then counterflip that bit to correct the error. You can do all this without ever determining what value each code bit holds. A similar strategy works to correct errors in a quantum system.

Learning the values of the parity checks still requires quantum measurement, but importantly, it does not reveal the underlying quantum information. Additional qubits can be used as disposable resources to obtain the parity values without revealing (and thus without disturbing) the encoded information itself.

What about no-cloning? It turns out it is possible to take a qubit whose state is unknown and encode that hidden state in a superposition across multiple qubits in a way that does not clone the original information. This process allows you to record what amounts to a single logical qubit of information across three physical qubits, and you can perform parity checks and corrective steps to protect the logical qubit against noise.

Quantum errors consist of more than just bit-flip errors, though, making this simple three-qubit repetition code unsuitable for protecting against all possible quantum errors. True QEC requires something more. That came in the mid-1990s when Peter Shor (then at AT&T Bell Laboratories, in Murray Hill, N.J.) described an elegant scheme to encode one logical qubit into nine physical qubits by embedding a repetition code inside another code. Shor's scheme protects against an arbitrary quantum error on any one of the physical qubits.

Since then, the QEC community has developed many improved encoding schemes, which use fewer physical qubits per logical qubit—the most compact use five—or enjoy other performance enhancements. Today, the workhorse of large-scale proposals for error correction in quantum computers is called the surface code, developed in the late 1990s by borrowing exotic mathematics from topology and high-energy physics.

**IT IS CONVENIENT** to think of a quantum computer as being made up of logical qubits and logical gates that sit atop an underlying foundation of physical devices. These physical devices are subject to noise, which creates physical errors that accumulate over time. Periodically, generalized parity measurements (called syndrome measurements) identify the physical errors, and corrections remove them before they cause damage at the logical level.

A quantum computation with QEC then consists of cycles of gates acting on qubits, syndrome measurements, error inference, and corrections. In terms more familiar to engineers, QEC is a form of feedback stabilization that uses indirect measurements to gain just the information needed to correct errors.

QEC is not foolproof, of course. The three-bit repetition code, for example, fails if more than one bit has been flipped. What's more, the resources and mechanisms that create the encoded quantum states and perform the syndrome measurements are themselves prone to errors. How, then, can a quantum computer perform QEC when all these processes are themselves faulty?

Remarkably, the error-correction cycle can be designed to tolerate errors and faults that occur at every stage, whether in the physical qubits, the physical gates, or even in the very measurements used to infer the existence of errors! Called a fault-tolerant architecture, such a design permits, in principle, error-robust quantum processing even when all the component parts are unreliable.

Even in a fault-tolerant architecture, the additional complexity introduces new avenues for failure. The effect of errors is therefore reduced at the logical level only if the underlying physical error rate is not too high. The maximum physical error rate that a specific fault-tolerant architecture can reliably handle is known as its break-even error threshold. If error rates are lower than this threshold, the QEC process tends to suppress errors over the entire cycle. But if error rates exceed the threshold, the added machinery just makes things worse overall.

The theory of fault-tolerant QEC is foundational to every effort to build useful quantum computers because it paves the way to building systems of any size. If QEC is implemented effectively on hardware exceeding certain performance requirements, the effect of errors can be reduced to arbitrarily low levels, enabling the execution of arbitrarily long computations.

At this point, you may be wondering how QEC has evaded the problem of continuous errors, which is fatal for scaling up analog computers. The answer lies in the nature of quantum measurements.

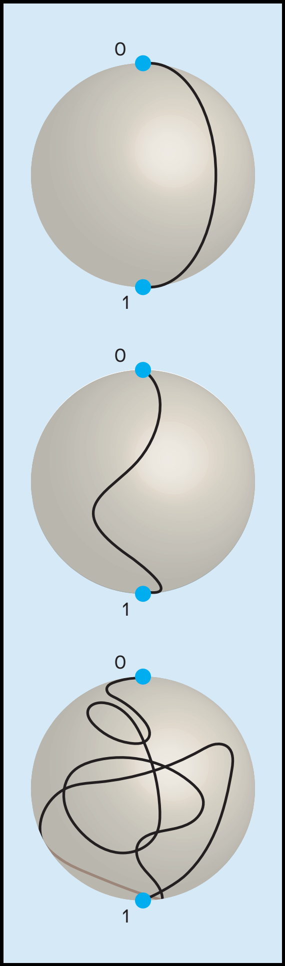

A superconducting qubit can be flipped by applying a simple microwave pulse that takes the qubit's state on a direct path on the Bloch sphere from 0 to 1 [top], but noise will introduce an error in the final position. A complicated pulse producing a more circuitous route can reduce the average amount of error in the final position. Here, the paths are chosen to minimize the effect of noise in the pulse amplitude alone [middle] or in both the amplitude and phase of the pulse [bottom].

In a typical quantum measurement of a superposition, only a few discrete outcomes are possible, and the physical state changes to match the result that the measurement finds. With the parity-check measurements, this change helps.

Imagine you have a code block of three physical qubits, and one of these qubit states has wandered a little from its ideal state (see diagram on page 31). If you perform a parity measurement, just two results are possible: Most often, the measurement will report the parity state that corresponds to no error, and after the measurement, all three qubits will be in the correct state, whatever it is. Occasionally the measurement will instead indicate the odd parity state, which means an errant qubit is now fully flipped. If so, you can flip that qubit back to restore the desired encoded logical state.

In other words, performing QEC transforms small, continuous errors into infrequent but discrete errors, similar to the errors that arise in digital computers.

**RESEARCHERS HAVE NOW** demonstrated many of the principles of QEC in the laboratory—from the basics of the repetition code through to complex encodings, logical operations on code words, and repeated cycles of measurement and correction. Current estimates of the break-even threshold for quantum hardware place it at about 1 error in 1,000 operations. This level of performance hasn't yet been achieved across all the constituent parts of a QEC scheme, but researchers are getting ever closer, achieving multiqubit logic with rates of fewer than about 5 errors per 1,000 operations. Even so, passing that critical milestone will be the beginning of the story, not the end.

On a system with a physical error rate just below the threshold, QEC would require enormous redundancy to push the logical rate down very far. It becomes much less challenging with a physical rate further below the threshold. So just crossing the error threshold is not sufficient—we need to beat it by a wide margin. How can that be done?

If we take a step back, we can see that the challenge of dealing with errors in quantum computers is one of stabilizing a dynamic system against external disturbances. Although the mathematical rules differ for the quantum system, this is a familiar problem in the discipline of control engineering. And just as control theory can help engineers build robots capable of righting themselves when they stumble, quantum-control engineering can suggest the best ways to implement abstract QEC codes on real physical hardware. Quantum control

CONTINUED ON PAGE 46

BY MITCHELL LAZARUS

# RADIO-SPECTRUM TURF WARS

## AND HOW THE FCC TRIES TO KEEP THE PEACE

**You've no doubt** seen the scary headlines: *Will 5G Cause Planes to Crash*? They appeared late last year, after the U.S. Federal Aviation Administration warned that new 5G services from AT&T and Verizon might interfere with the radar altimeters that airplane pilots rely on to land safely. Not true, said AT&T and Verizon, with the backing of the U.S. Federal Communications Commission, which had authorized 5G. The altimeters are safe, they maintained. Air travelers didn't know what to believe.

Another recent FCC decision had also created a controversy about public safety: okaying Wi-Fi devices in a 6-gigahertz frequency band long used by point-to-point microwave systems to carry safety-critical data. The microwave operators predicted that the Wi-Fi devices would disrupt their systems; the Wi-Fi interests insisted they would not. (As an attorney, I represented a microwave-industry group in the ensuing legal dispute.)

Whether a new radio-based service will interfere with existing services in the same slice of the spectrum seems like a straightforward physics problem. Usually, though, opposing parties' technical analyses give different results. Disagreement among the engineers then opens the way for public safety to become just one among several competing interests. I've been in the thick of such arguments, so I wanted to share how these issues arise and how they are settled.

**Not all radio spectrum** is created equal. Lower frequencies travel farther and propagate better through buildings and terrain. Higher frequencies offer the bandwidth to carry more data, and work well with smaller antennas. Every radio-based application has its own needs and its own spectral sweet spot.

Suitable spectrum for mobile data—4G, 5G, Wi-Fi, Bluetooth, many others—runs from a few hundred megahertz to a few gigahertz. Phones, tablets, laptops, smart speakers, Wi-Fi-enabled TVs and other appliances, Internet-of-things devices, lots of commercial and industrial gear—they all need these same frequencies.

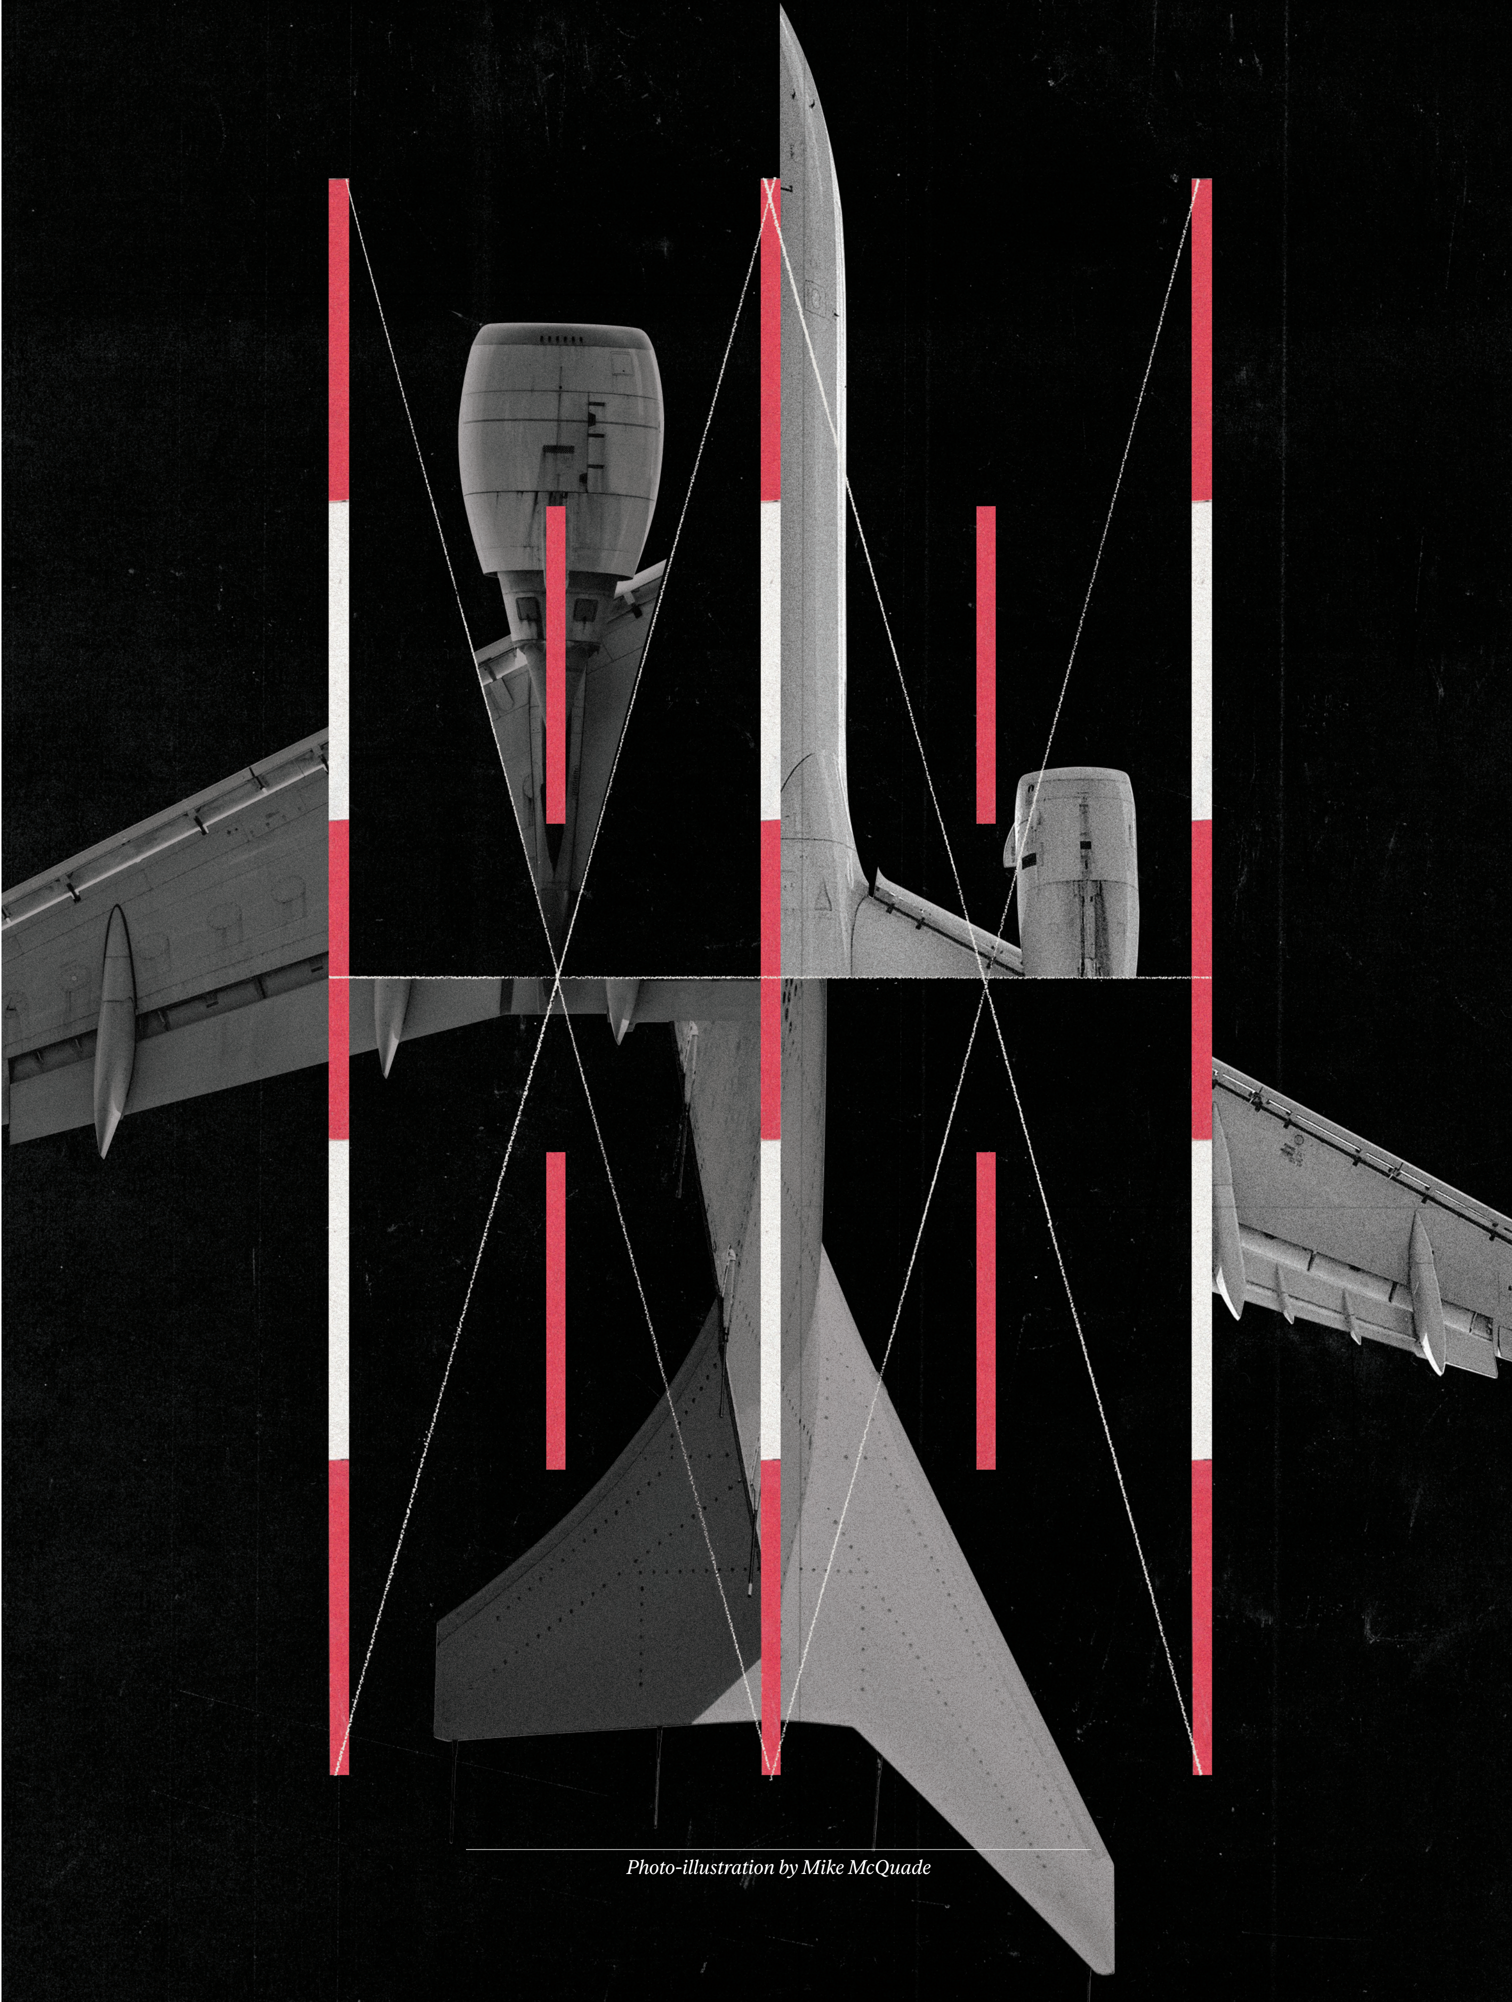

*Photo-illustration by Mike McQuade*

The problem is that this region of spectrum has been fully occupied for decades. So when a new service like 5G appears, or an older one like Wi-Fi needs room to expand, the FCC has two options. For a licensed service like 5G, the FCC generally clears incumbent users from a range of frequencies—either repacking them into other frequencies nearby or relocating them to a different part of the spectrum—and then auctions the freed-up spectrum to providers of the new service. To accommodate an unlicensed service like Wi-Fi, the FCC overlays the new users onto the same frequencies as the incumbents, usually at lower power.

The FCC tries to write technical rules for the new or expanded service that will leave the incumbents mostly unaffected. It is commonplace for newcomers to complain that any interference they cause is not their fault, attributing it to inferior incumbent receivers that fail to screen out unwanted signals. This argument usually fails. The newcomer must deal with the spectrum and its occupants as it finds them. Strategies for accomplishing that task vary.

**Congress prohibits the FCC** (and other federal agencies) from changing the regulatory ground rules without first soliciting and considering public input. On technical issues, that input comes mostly from the affected industries after the FCC outlines its tentative plans in a Notice of Proposed Rulemaking. There follows a back-and-forth exchange of written submissions posted to the FCC's website, typically lasting a year or more.

Ordinarily, parties can also make in-person presentations to the FCC staff and the five commissioners, if they post summaries of what they say. Sometimes the staff uses these meetings to test possible compromises among the parties.

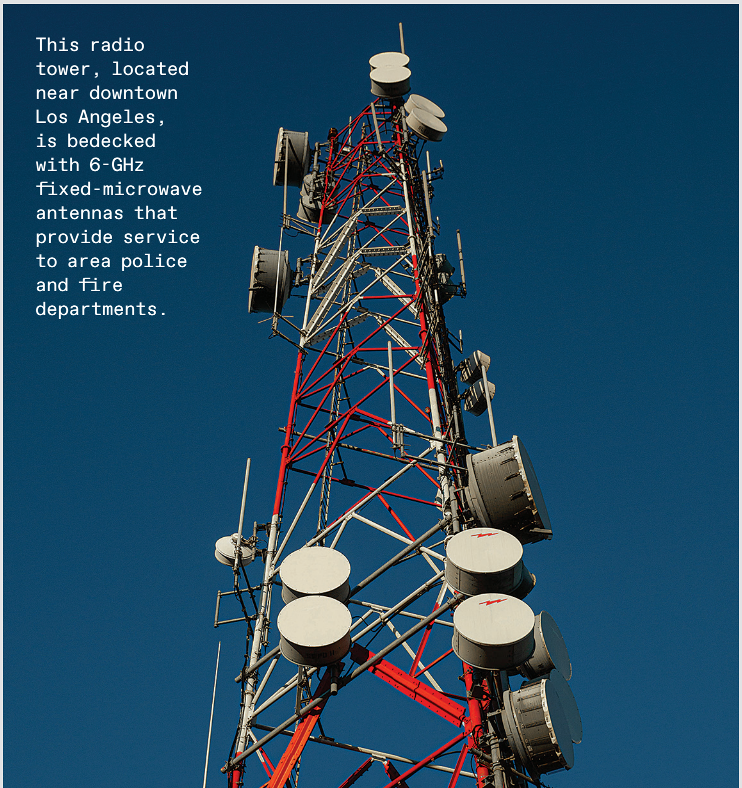
This radio tower, located near downtown Los Angeles, is bedecked with 6-GHz fixed-microwave antennas that provide service to area police and fire departments.

GEORGE ROSE/GETTY IMAGES

All this openness and transparency has a big exception: Other federal agencies, like the FAA, can and sometimes do submit comments to the FCC's website, but they also have a back channel to deliver private communications.

The submissions in a spectrum proceeding generally make two kinds of points. First, the newcomers and the incumbents both present data to impress the FCC with their respective services' widespread demand, importance to the economy, and utility in promoting education, safety, and other public benefits. Second, both the proponents and opponents of a new frequency usage submit engineering studies and simulations, sometimes running to hundreds of pages.

Predictably, the two parties' studies come to opposite conclusions. The proponents show the new operations will have no harmful effect on incumbents, while the incumbents demonstrate that they will suffer devastating interference. Each party responds with point-by-point critiques of the other side's studies and may carry out counterstudies for further proof the other side is wrong.

How do such alternative realities arise? It's not because they are based on different versions of Maxwell's equations. The two sides' studies usually disagree because they start with differing assumptions about the newcomer's transmitter characteristics, the incumbent's receiver characteristics, and the geometries and propagation that govern interaction between the two. Small changes to some of these factors can produce large changes in the results.

Sometimes the parties, the FCC, or another government agency may conduct hardware tests in the lab or in the field to assess the degree of interference and its effects. Rather than settle anything, though, these experiments just add fuel to the controversy. Parties disagree on whether the test setup was realistic, whether the data were analyzed correctly, and what the results imply for real-world operations.

When, for example, aviation interests ran tests that found 5G transmissions caused interference to radio altimeters, wireless carriers vigorously challenged their results. In contrast, there was no testing in the 6-GHz Wi-Fi proceeding, where the disagreements turned on theoretical analyses and simulations.

Further complicating matters, the disputed studies and tests do not predict interference as a binary yes/no but as differing probabilities for various degrees of interference. And the parties involved often disagree on whether a given level of interference is harmless or will cause the victim's receiver to malfunction. Reaching a decision on interference issues requires the FCC to make its way through a multidimensional maze of conflicting uncertainties. Here are some concrete issues that illuminate this all-too-common dynamic.

**Those ubiquitous sideways-facing** dishes on towers and buildings are fixed-microwave antennas. Equipment of this kind has operated reliably since

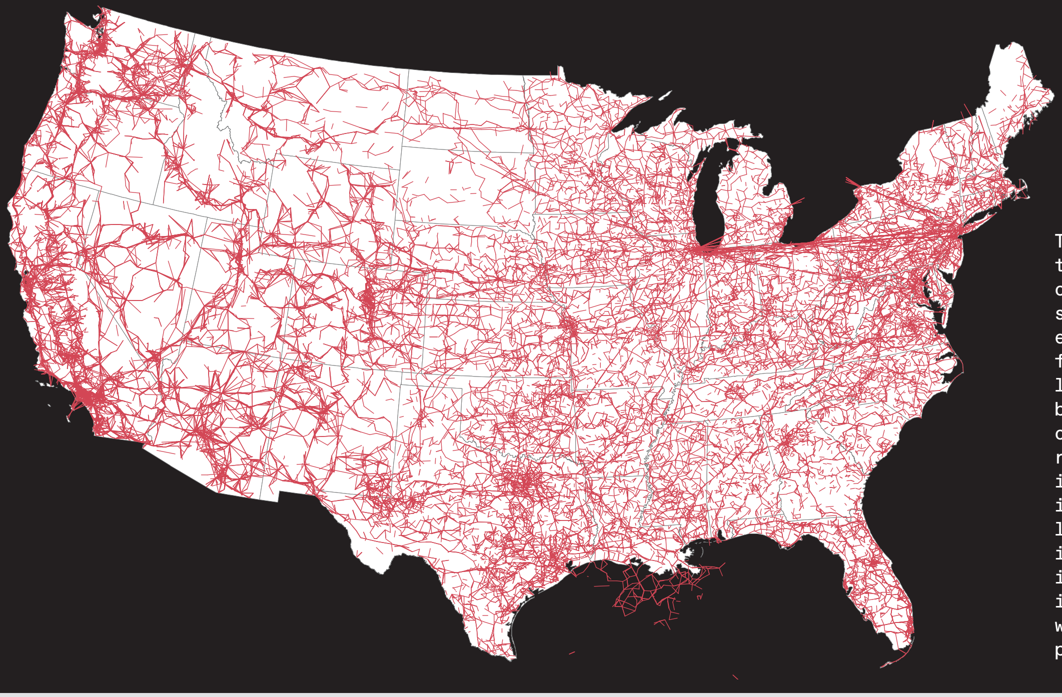

The red lines on this map of the 48 contiguous U.S. states show the location of existing 6-gigahertz fixed-microwave links, as recorded by Comsearch, a company that helps relevant parties avoid issues with radio interference. These links connect people in almost all areas, including far offshore in the Gulf of Mexico, where drilling platforms are common.

COMSEARCH

the 1950s. The 6-GHz band, the lowest-frequency microwave band available today, is the only one capable of 100-kilometer hops, making it indispensable. Along with more pedestrian uses, the band carries safety-critical information: to coordinate trains, control pressure in oil and gas pipelines, balance the electric grid, manage water utilities, and route emergency telephone calls.

Four years ago, when the FCC proposed adding Wi-Fi to the 6-GHz band, all sides agreed that the vast majority of Wi-Fi devices would cause no trouble. Statistically, most would be outside the microwave antennas' highly directional main beams, or on the wrong frequency, or shielded by buildings, terrain, and ground clutter.

The dispute centered on the small proportion of devices that might transmit on a frequency in use while being in the line-of-sight of a microwave antenna. The Wi-Fi proponents projected just under a billion devices, operating among 100,000 microwave receivers. The opponents pointed out that even a very small fraction of the many new transmitters could cause troubling numbers of interference events.

To mitigate the problem, the FCC adopted rules for an Automatic Frequency Control (AFC) system. A Wi-Fi device must either report its location to a central AFC database, which assigns it noninterfering frequencies for that location, or operate close to and under the control of an AFC-guided device. The AFC system will not be fully operational for another year or two, and disagreements persist about the details of its eventual operation.

More controversially, the FCC also authorized Wi-Fi devices without AFC, transmitting at will on any 6-GHz frequency from any geographic location—but only indoors and at no more than one-quarter of the maximum AFC-controlled power. The Wi-Fi proponents' technical studies showed that attenuation from building walls would prevent interference. The microwave operators' studies showed the opposite: that interference from uncontrolled indoor devices was virtually certain.

How could engineers, using the same equations, come to such different conclusions? These are a few of the ways in which their analyses differed:

**Wi-Fi device power:** A Wi-Fi device transmits in short bursts, active about 1/250th of the time, on average. The Wi-Fi proponents scaled down the power by a like amount, treating a device that transmits intermittently at, say, 250 milliwatts as though it transmitted continuously at 1 mW. The microwave operators argued that interference can occur only while the device is actually transmitting, so they calculated using the full power.

**Building attenuation:** A 6-GHz signal encounters substantial attenuation from concrete building walls and thermal windows, less from wood walls, and practically none from plain-glass windows. The Wi-Fi proponents took weighted averages over several building materials to calculate typical wall attenuations. The microwave operators reasoned that interference was most likely from an atypical Wi-Fi device behind plain glass, and they calculated accordingly, assuming a minimal amount of attenuation.

**Path loss:** In estimating the signal loss from a building that houses a Wi-Fi device to a microwave-receiving antenna, the Wi-Fi proponents used a standard propagation model that incorporates attenuation due to other buildings, ground clutter, and the like. The microwave operators were most concerned about a device located with open air between the building and the antenna, so they used free-space propagation in their calculations.

Using their preferred starting assumptions, the Wi-Fi proponents proved that Wi-Fi devices over a wide range of typical situations present no risk of interference. Using a different set of assumptions, the microwave operators proved there is a large risk of interference from a small proportion of Wi-Fi devices in atypical locations, arguing that multiplying that small proportion by almost a billion Wi-Fi devices made interference virtually certain.

**Americans want their smartphones** and tablets to have fast Internet access everywhere. That takes a lot of spectrum. Congress passed a statute in 2018 that told the FCC to find more—and specifically to consider 3.7 to 4.2 GHz, part of the C-band, used since the 1960s to receive satellite signals. The FCC partitioned the band in 2020, allocating 3.7 to 3.98 GHz for 5G mobile data. In early 2021, it auctioned the new 5G frequencies for US $81 billion, mostly to Verizon and AT&T. The auction winners were also expected to pay the satellite providers around $13 billion to compensate them for the costs of moving to other frequencies.

A nearby band at 4.2 to 4.4 GHz serves radar altimeters (also called radio altimeters), instruments that tell a pilot or an automatic landing system how high the aircraft is above the ground. The altimeter works by emitting downward radio waves that reflect off the ground and back up to a receiver in the device. The time for the round trip gives the altitude. Large planes operate two or three altimeters simultaneously, for redundancy.

Even though the altimeters use frequencies separated from the 5G band, they can still receive interference from 5G. That's because every transmitter, including ones used for 5G, emits unwanted signals outside its assigned frequencies. Every receiver is likewise sensitive to signals outside its intended range, some more than others. Interference can occur if energy from a 5G transmitter falls within the sensitivity range of the receiver in an altimeter.

The FCC regulates transmitter out-of-band emissions. In contrast, it has few rules on receiver out-of-band reception (although it recently opened a discussion on whether to expand them). Manufacturers generally design receivers to function reliably in their expected environments, which can leave them vulnerable if a new service appears in formerly quiet spectrum near the frequencies they receive on.

Aviation interests feared this outcome with the launch of C-band 5G, one citing the possibility of "catastrophic impact with the ground, leading to multiple fatalities." The FCC's 5G order tersely dismissed concerns about altimeter interference, although it invited the aviation industry to study the matter further. The industry did so, renewing its concerns and requesting that the wireless carriers refrain from using 5G near airports. But this came *after* the wireless carriers had committed almost $100 billion and begun building out facilities.

Much as in the case of 6-GHz Wi-Fi, the 5G providers and aviation advocates reached different predictions about interference by starting with different assumptions. Some key areas of disagreement were:

5G out-of-band emissions: The radio engineers working for aviation groups assumed higher levels than the wireless carriers, which said the numbers in the aviation study levels exceeded FCC limits.

Off-channel sensitivity in altimeter receivers: There are several makes and models of altimeters in use, having varying receiver characteristics, leading to disagreements on which to include in the studies.

Altimeters in the same or other aircraft nearby: A busy airport has a lot of altimeters operating. Wireless carriers said these would overpower 5G interference. Aviation interests countered that multiple altimeters in the area would consume one another's interference margin and leave them all more vulnerable to 5G.

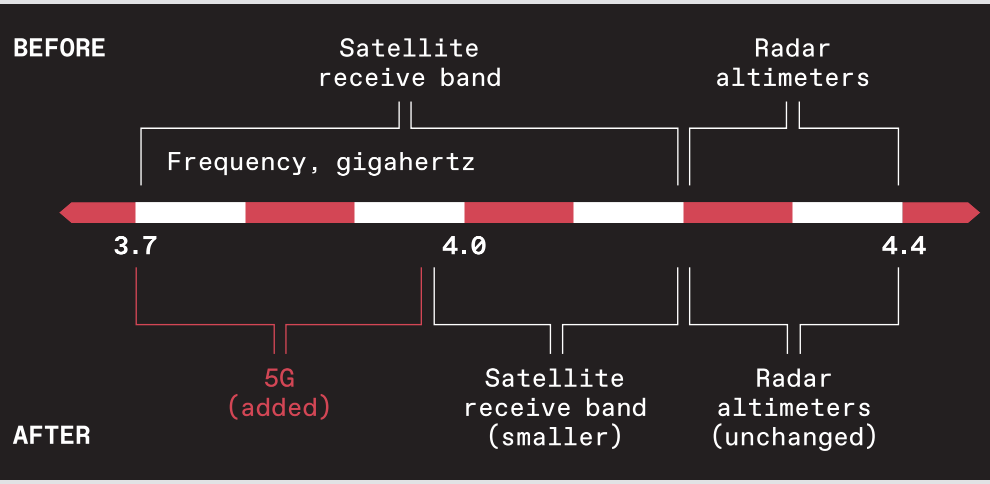

To make way for new 5G cellular services, the Federal Communications Commission reallocated part of the radio spectrum. That reallocation resulted in 5G transmissions that are close in frequency to a band used by aircraft radar altimeters.

These radio towers, which sit atop Black Mountain in Carmel Valley, Calif., include many drumlike antennas used for 6-gigahertz fixed-microwave links.

SHUTTERSTOCK

**Aircraft pitch and roll:** Aviation advocates argued that the changing angles of the aircraft as it approaches the runway can expose the altimeter receivers to more 5G signal.

**Reflectivity of the ground:** Aviation groups favored modeling with lower values of reflectivity, which reduce the received signal strength at the altimeter and hence increase its susceptibility to 5G interference.

The carriers temporarily paused 5G rollout near some airports, and the airlines canceled and rescheduled some flights. At this writing, the FAA is evaluating potentially affected aircraft, altimeters, and airport systems. Most likely, 5G will prevail. In the extremely improbable event that the FAA and the FCC were to agree that C-band 5G cannot operate safely near airports, the wireless carriers presumably would be entitled to a partial refund of their $81 billion auction payments.

**Making complicated trade-offs** has long been the job of the five FCC commissioners. They are political appointees, nominated by the president and confirmed by the Senate. The four now in office (there is a vacancy) are all lawyers. It has been decades since a commissioner had a technical background. The FCC has highly capable engineers on staff, but only in advisory roles. The commissioners have no obligation to take their advice.

Congress requires the FCC to regulate "in the public interest," but the commissioners must determine what that means in each case. Legally, they can reach any result that has at least some support in the submissions, even if other submissions more strongly support an opposite result. Submissions to the FCC in both the 6-GHz and 5G matters conveyed sharp disagreement as to how much safety protection the public interest requires.

To fully protect 6-GHz microwave operations against interference from the small fraction of Wi-Fi devices in the line of sight of the microwave receivers would require degrading Wi-Fi service for large numbers of people. Similarly, eliminating any chance whatsoever of a catastrophic altimeter malfunction due to 5G interference might require turning off C-band 5G in some heavily populated areas.

The orders that authorized 6-GHz Wi-Fi and C-band 5G did not go that far and did not claim they had achieved zero risk. The order on 5G stated that altimeters had "all due protection." In the 6-GHz case, with a federal appeals court deferring to its technical expertise, the FCC said it had "reduce[d] the possibility of harmful interference to the minimum that the public interest requires."

These formulations make clear that safety is just one of several elements in the mix of public interests considered. Commissioners have to balance the goals of minimizing the risk of plane crashes and pipeline explosions against the demand for ubiquitous Internet access and Congress's mandate to repurpose more spectrum.

In the end, the commissioners agreed with proponents' claims that the risk of harmful interference from 6-GHz Wi-Fi is "insignificant," although not zero, and similarly from 5G, not "likely...under... reasonably foreseeable scenarios"—conclusions that made it possible to offer the new services.

People like to think that the government puts the absolute safety of its citizens above all else. Regulation, though, like engineering, is an ever-shifting sequence of trade-offs. The officials who set highway speed limits know that lower numbers will save lives, but they also take into account motorists' wishes to get to their destinations in a timely way. So it shouldn't come as a great surprise that the FCC performs a similar balancing act. ■

# WEANING INDIA FROM COAL

By ASHOK JHUNJHUNWALA,
KAUSHAL KUMAR JHA
& ANSON SANDO

## How to hasten the country's transition to solar and wind power

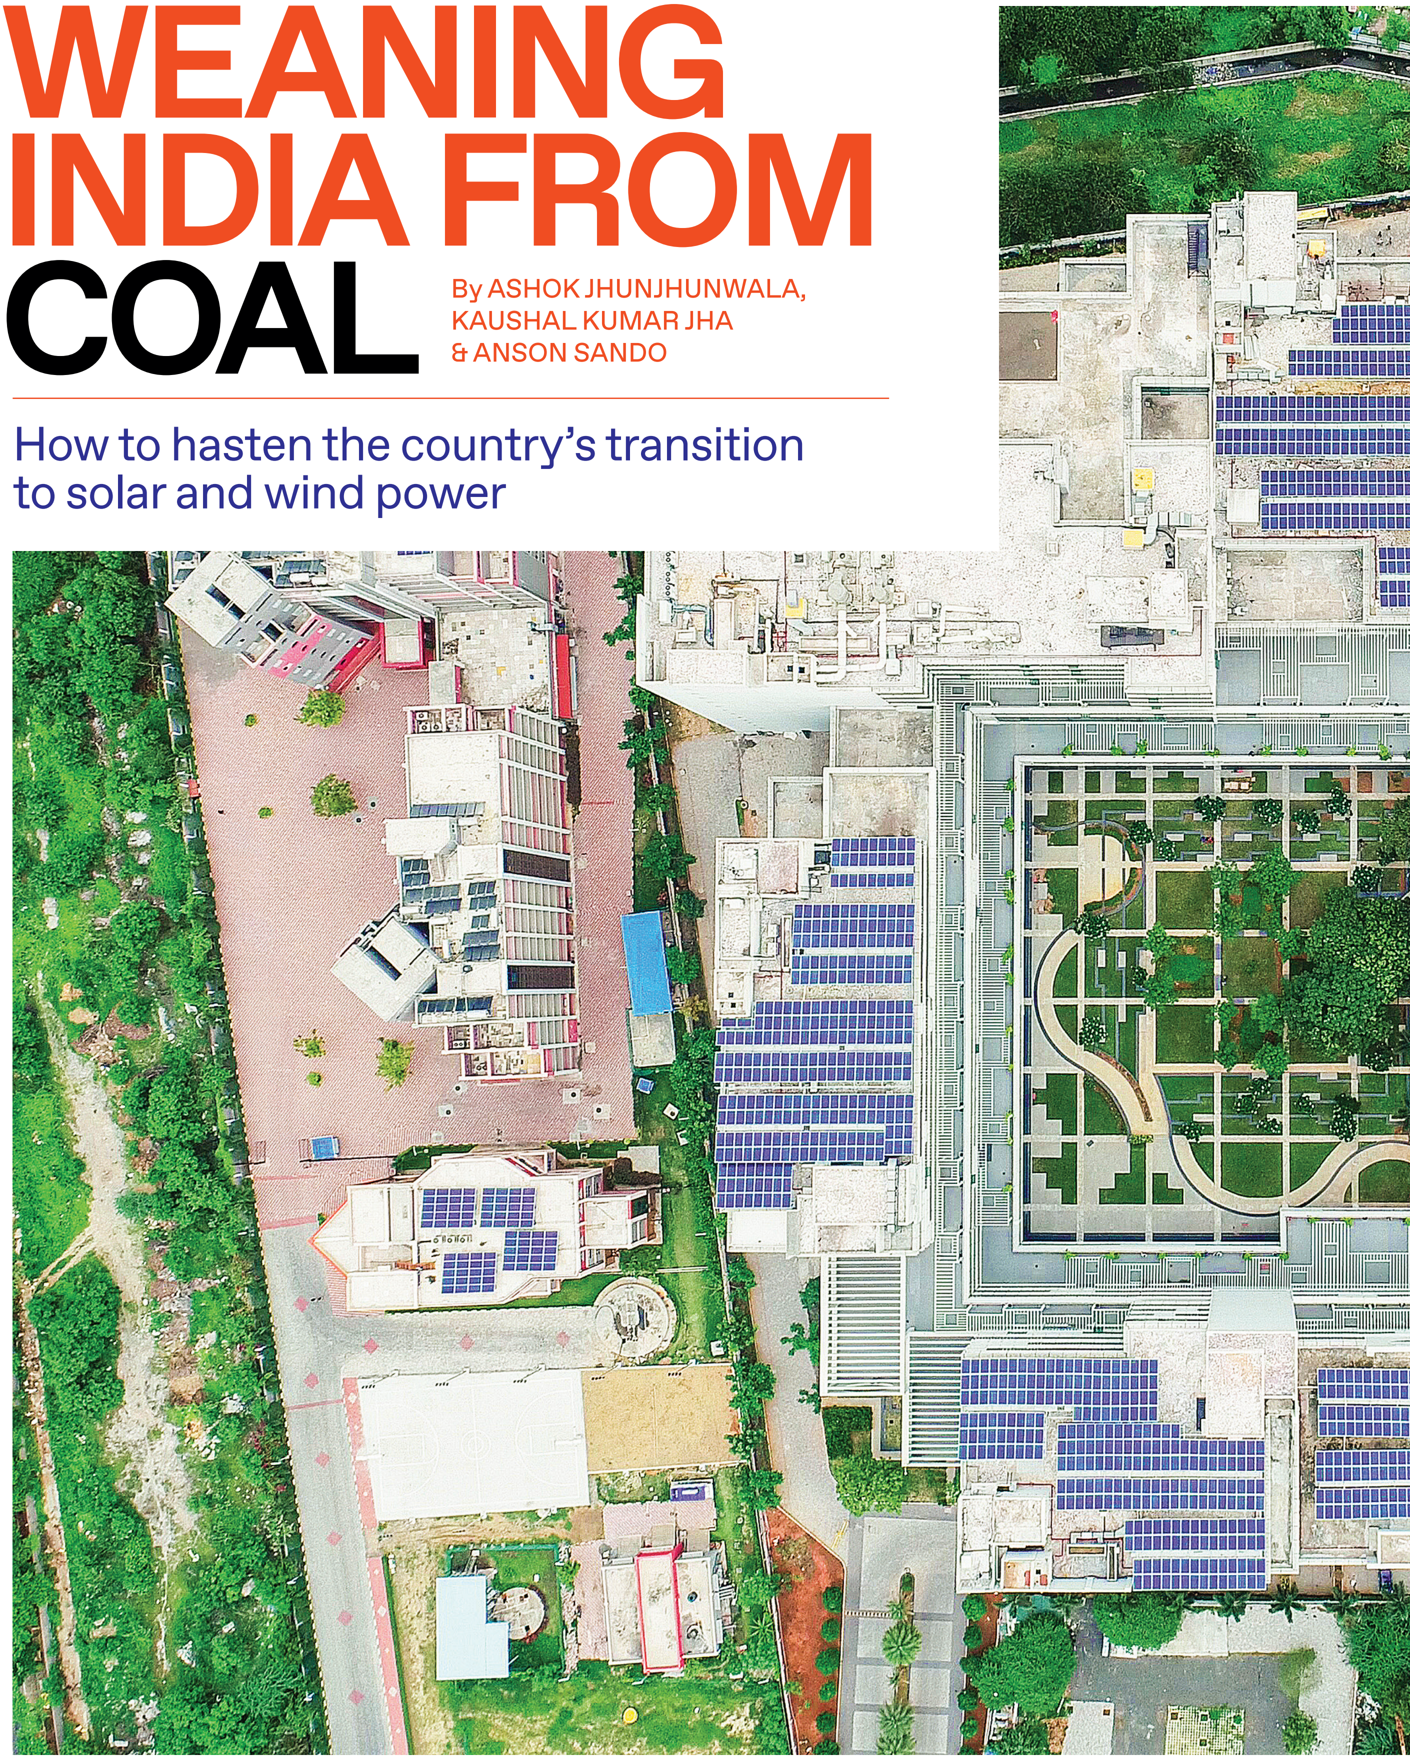

The IITM Research Park, in Chennai, provides R&D facilities for hundreds of companies. Rooftop solar provides about 10 percent of the complex's electricity. The addition of dedicated off-site solar and wind power plus on-site energy storage should allow IITMRP's renewable energy usage to move closer to 100 percent in the next few years.

**THE RISING THREAT** of global warming requires that every country act now. The question is how much any one country should do.

India is 126th in the world in per capita carbon dioxide emissions, according to a 2020 European Union report. One might argue that the onus of reversing global warming should fall on the developed world, which on a per-capita basis consumes much more energy and emits significantly more greenhouse gases. However, India ranks third in the world in total greenhouse gas emissions—the result of having the second-largest population and being third largest in energy consumption.

As India's GDP and per capita income continue to climb, so too will its energy consumption. For instance, just 8 percent of Indian homes had air-conditioning in 2018, but that share is likely to rise to 50 percent by 2050. The country's electricity consumption in 2019 was nearly six times as great as in 1990. Greenhouse gas emissions will certainly grow too, because India's energy generation is dominated by fossil fuels—coal-fired power plants for electricity, coal- and gas-fired furnaces for industrial heating, liquid petroleum gas for cooking, and gasoline and diesel for transportation.

Fossil fuels dominate even though renewable energy generation in many parts of the world now costs less than fossil-fuel-based electricity. While electricity from older coal plants in India costs 2.7 U.S. cents per kilowatt-hour and 5.5 cents from newer plants that have additional pollution-control equipment, the cost of solar energy has dropped to 2.7 cents per kilowatt-hour, and wind power to 3.4 cents per kilowatt-hour. As renewable energy has steadily gotten cheaper, the installed capacity has grown, to 110 gigawatts. That amounts to 27 percent of capacity, compared to coal's share, which is 52 percent. The government of India has set a target of 450 GW of renewable energy capacity by 2030.

Yet in terms of energy *generated*, renewable energy in India still falls short. In 2021, about 73 percent of the country's electricity was produced from coal, and only 9.6 percent from solar and wind power. That's because solar and wind power aren't available around the clock, so the proportion of the installed capacity that gets used is just 20 to 30 percent. For coal, the capacity utilization rate can go as high as 90 percent.

As renewable energy capacity grows, the only way to drastically reduce coal in the electricity mix is by adding energy storage. Although some of the newer solar plants and wind farms are being set up with large amounts of battery storage, it could be decades before such investments have a significant impact. But there is another way for India to move faster toward its decarbonization goal: by focusing the renewable-energy push in India's commercial and industrial sectors.

India has some 40,000 commercial complexes, which house offices and research centers as well as shopping centers and restaurants. Together they consume about 8 percent of the country's electricity. The total footprint of such complexes is expected to triple by 2030, compared to 2010. To attract tenants, the managers of these complexes like to project their properties as users of renewable energy.

India's industrial sector, meanwhile, consumes about 40 percent of the country's electricity, and many industrial operators would also be happy to adopt a greater share of renewable energy if they can see a clear return on investment.

A 2-megawatt solar plant located about 500 kilometers from IITM Research Park provides dedicated electricity to the complex. A 2.1-MW wind farm now under construction will feed IITMRP through a similar arrangement.

ALL PHOTOS: IIT MADRAS

The IITMRP's chilled-water system provides air-conditioning to the complex. Water is chilled to about 6 °C and then stored in this 300-cubic-meter underground tank for later circulation to the offices.

Right now, many of these complexes use rooftop solar, but limited space means they can only get a small share of their energy that way. These same complexes can, however, leverage a special power-transmission and "wheeling" policy that's offered in India. Under this arrangement, an independent power-generation company sets up solar- or wind-power plants for multiple customers, with each customer investing in the amount of capacity it needs. In India, this approach is known as a group-captive model. The generating station injects the electricity onto the grid, and the same amount is immediately delivered, or wheeled in, to the customer, using the utility's existing transmission and distribution network. A complex can add energy storage to save any excess electricity for later use. If enough commercial, industrial, and residential complexes adopt this approach, India could rapidly move away from coal-based electricity and meet a greater share of its energy needs with renewable energy. Our group at the Indian Institute of Technology Madras has been developing a pilot to showcase how a commercial complex can benefit from this approach.

**THE COMMERCIAL COMPLEX** known as the IITM Research Park, or IITMRP, in Chennai, is a 110,000-square-meter facility that houses R&D facilities for more than 250 companies, including about 150 startups, and employs about 5,000 workers. It uses an average of 40 megawatt-hours of electricity per weekday, or about 12 gigawatt-hours per year. Within the campus, there is 1 megawatt of rooftop solar, which provides about 10 percent of IITMRP's energy. The complex is also investing in 2 MW of captive solar and 2.1 MW of captive wind power off-site, the electricity from which will be wheeled in. This will boost the renewable-energy usage to nearly 90 percent in about three years. Should the local power grid fail, the complex has backup diesel generators.

Of course, the generation of solar and wind energy varies from minute to minute, day to day, and season to season. The total generated energy will rarely meet IITMRP's demand exactly; it will usually either exceed demand or fall short.

To get closer to 100 percent renewable energy, the complex needs to store some of its wind and solar power. To that end, the complex is building two complementary kinds of energy storage. The first is a 2-MWh, 750-volt direct-current lithium-ion battery facility. The second is a chilled-water storage system with a capacity equivalent to about 2.45 MWh. Both systems were designed and fabricated at IITMRP.

The battery system's stored electricity can be used wherever it's needed. The chilled-water system serves a specific, yet crucial function: It helps cool the buildings. For commercial complexes in tropical climates like Chennai's, nearly 40 percent of the energy goes toward air-conditioning, which can be costly. In the IITMRP system, a central heating, ventilation, and air-conditioning (HVAC) system chills water to about 6 °C, which is then circulated to each office. A 300-cubic-meter underground tank stores the chilled water for use within about 6 to 8 hours. That relatively short duration is because the temperature of the chilled water in the tank rises about 1 °C every 2 hours.

The heat transfer capacity of the chilled-water system is 17,500 megajoules, which as mentioned is equivalent to 2.45 MWh of battery storage. The end-to-end round-trip energy loss is about 5 percent. And unlike with a battery system, you can “charge” and “discharge” the chilled-water tank several times a day without diminishing its life span.

**ALTHOUGH ENERGY STORAGE** adds to the complex’s capital costs, our calculations show that it ultimately reduces the cost of power. The off-site solar and wind farms are located, respectively, 500 and 600 kilometers from IITMRP. The cost of the power delivered to the complex includes generation (including transmission losses) of 5.14 cents/kWh as well as transmission and distribution charges of 0.89 cents/kWh. In addition, the utilities that supply the solar and wind power impose a charge to cover electricity drawn during times of peak demand. On average, this demand charge is about 1.37 cents/kWh. Thus, the total generation cost for the solar and wind power delivered to IITMRP is about 7.4 cents/kWh.

There’s also a cost associated with energy storage. Because most of the renewable energy coming into the complex will be used immediately, only the excess needs to be stored—about 30 percent of the total, according to our estimate.

So the average cost of round-the-clock renewable energy works out to 9.3 cents/kWh, taking into account the depreciation, financing, and operation costs over the lifetime of the storage. In the future, as the cost of energy storage continues to decline, the average cost will remain close to 9 cents/kWh, even if half of the energy generated goes to storage. And the total energy cost could drop further with declines in interest rates, the cost of solar and wind energy, or transmission and demand charges.

For now, the rate of 9.3 cents/kWh compares quite favorably to what IITMRP pays for regular grid power—about 15 cents/kWh. That means, with careful design, the complex can approach 100 percent renewable energy and still save about a third on the energy costs that it pays today. Keep in mind that grid power in India primarily comes from coal-based generation, so for IITMRP and other commercial complexes, using renewable energy plus storage has a big environmental upside.

Electricity tariffs are lower for India’s industrial and residential sectors, so the cost advantage of this approach may not be as pronounced in those settings. But renewable energy can also be a selling point for the owners of such complexes—they know many tenants like having their business or home located in a property that’s green.

**ALTHOUGH IITMRP’S** annual consumption is about 12 GWh, the energy usage, or load, varies slightly from month to month, from 970 to 1,100 MWh. Meanwhile, the energy generated from the captive off-site solar and wind plants and the rooftop solar plant will vary quite a bit more. The top chart on this page shows the estimated monthly energy generated and the monthly load.

As is apparent, there is some excess energy available in May and July, and an overall energy deficit at other times. In October, November, and December, the deficit is substantial, because wind-power generation tends to be lowest during those months. Averaged over a year, the deficit works out to be 11 percent; the arrangement we’ve described, in other words, will allow IITMRP to obtain 89 percent of its energy from renewables.

For the complex to reach 100 percent renewable energy, it’s imperative that any excess energy be stored and then used later to make up for the renewable energy deficits. When the energy deficits are particularly high, the only way to boost renewable energy usage further will be to add another source of generation, or else add long-term

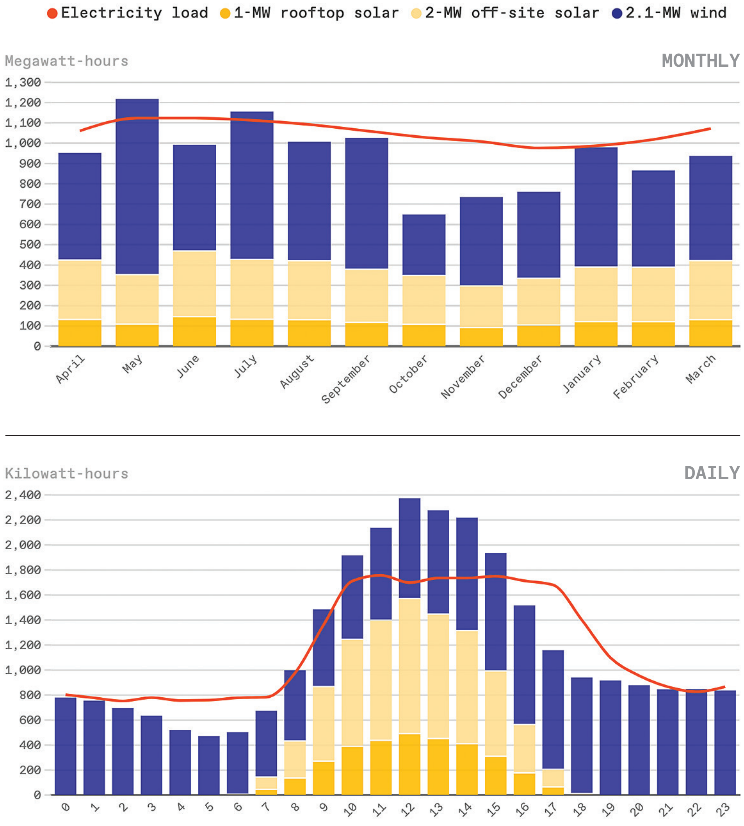

IITM Research Park’s electricity load and available renewable energy vary across months [top] and over the course of a single day [bottom]. To make up for the deficit in renewable energy, especially during October, November, and December, additional renewable generation or long-term energy storage will be needed. At other times of the year, the available renewable energy tends to track the load closely throughout the day, with any excess energy sent to storage. SOURCE: IIT MADRAS

IITM Research Park's lithium-ion battery facility stores excess electricity for use after the sun goes down or when there's a dip in wind power.

energy storage that's capable of storing energy over months. Researchers at IITMRP are working on additional sources of renewable energy generation, including ocean, wave, and tidal energy, along with long-term energy storage, such as zinc-air batteries.

For other times of the year, the complex can get by on a smaller amount of shorter-term storage. How much storage? If we look at the energy generated and the load on an hourly basis over a typical weekday, we see that the total daily load generally matches the total daily demand, but with small fluctuations in surplus and deficit. Those fluctuations represent the amount of energy that has to move in and out of storage. In the bottom chart on page 44, the cumulative deficit peaks at 1.15 MWh, and the surplus peaks at 1.47 MWh. Thus, for much of the year, a storage size of 2.62 MWh should ensure that no energy is wasted.

This is a surprisingly modest amount of storage for a complex as large as IITMRP. It's possible because for much of the year, the load follows a pattern similar to the renewable energy generated. That is, the load peaks during the hours when the sun is out, so most of the solar energy is used directly, with a small amount of excess being stored for use after the sun goes down. The load drops during the evening and at night, when the wind power is enough to meet most of the complex's demand, with the surplus again going into storage to be used the next day, when demand picks up.

On weekends, the demand is, of course, much less, so more of the excess energy can be stored for later use on weekdays. Eventually, the complex's lithium-ion battery storage will be expanded to 5 MWh, to take advantage of that energy surplus. The batteries plus the chilled-water system will ensure that enough storage is available to take care of weekday deficits and surpluses most of the time.

As mentioned earlier, India has some 40,000 commercial complexes like IITMRP, and that number is expected to grow rapidly. Deploying energy storage for each complex and wheeling in solar and wind energy make sense both financially and environmentally. Meanwhile, as the cost of energy storage continues to fall, industrial complexes and large residential complexes could be enticed to adopt a similar approach. In a relatively short amount of time—a matter of years, rather than decades—renewable energy usage in India could rise to about 50 percent.

On the way to that admittedly ambitious goal, the country's power grids will also benefit from the decentralized energy management within these complexes. The complexes will generally meet their own supply and demand, enabling the grid to remain balanced. And, with thousands of complexes each deploying megawatts' worth of stationary batteries and chilled-water storage, the country's energy-storage industry will get a big boost. Given the government's commitment to expanding India's renewable capacity and usage, the approach we're piloting at IITMRP will help accelerate the push toward cleaner and greener power for all. ■

CONTINUED FROM PAGE 33

can minimize the effects of noise and make QEC practical.

In essence, quantum control involves optimizing how you implement all the physical processes used in QEC—from individual logic operations to the way measurements are performed. For example, in a system based on superconducting qubits, a qubit is flipped by irradiating it with a microwave pulse. One approach uses a simple type of pulse to move the qubit's state from one pole of the Bloch sphere, along the Greenwich meridian, to precisely the other pole. Errors arise if the pulse is distorted by noise. It turns out that a more complicated pulse, one that takes the qubit on a well-chosen meandering route from pole to pole, can result in less error in the qubit's final state under the same noise conditions, even when the new pulse is imperfectly implemented.

One facet of quantum-control engineering involves careful analysis and design of the best pulses for such tasks in a particular imperfect instance of a given system. It is a form of open-loop (measurement-free) control, which complements the closed-loop feedback control used in QEC.

This kind of open-loop control can also change the statistics of the physical-layer errors to better comport with the assumptions of QEC. For example, QEC performance is limited by the worst-case error within a logical block, and individual devices can vary a lot. Reducing that variability is very beneficial. In an experiment our team performed using IBM's publicly accessible machines, we showed that careful pulse optimization reduced the difference between the best-case and worst-case error in a small group of qubits by more than a factor of 10.

Some error processes arise only while carrying out complex algorithms. For instance, crosstalk errors occur on qubits only when their neighbors are being manipulated. Our team has shown that embedding quantum-control techniques into an algorithm can improve its overall success by orders of magnitude. This technique makes QEC protocols much more likely to correctly identify an error in a physical qubit.

For 25 years, QEC researchers have largely focused on mathematical strategies for encoding qubits and efficiently detecting errors in the encoded sets. Only recently have investigators begun to address the thorny question of how best to implement the full QEC feedback loop in real hardware. And while many areas of QEC technology are ripe for improvement, there is also growing awareness in the community that radical new approaches might be possible by marrying QEC and control theory. One way or another, this approach will turn quantum computing into a reality—and you can carve that in stone. ■

UNIVERSITY
of HAWAIʻI®
MĀNOA

**Assistant, Associate or Full Professor, position number 84464, University of Hawai'i at Mānoa (UHM), Department of Electrical and Computer Engineering** invites applications for a full-time, tenure-track faculty position, pending position clearance and availability of funds. To begin approximately **January 1, 2023** or soon thereafter.

**The University of Hawaiʻi at Mānoa**, a Carnegie R1 research university, is a top-50 public university dedicated to providing world-class teaching, research, and service in a multicultural and inclusive environment. Collaboration, funding opportunities, resources, and research exposure may be found through the department's involvement and affiliation with the $5 million **NSF CyberCorps Scholarship for Service Program (SFS)**; the **NSA/DHS National Center of Academic Excellence in Cyber Defense Research (CAE-R); the NSF AI Research Institute** (focusing on dynamic systems); **the NSF Center for Science of Information** (focusing on big data, information theory, and machine learning); and Ike Wai, a $20 million NSF EPSCOR project for bio/nano sensing. Further information is available at **https://ee.hawaii.edu**.

We are seeking candidates with a strong research record in microelectronics and related areas. Of particular interest are candidates in the areas of mixed-signal integrated circuit design, embedded machine learning, sensor interfaces, or related areas. However, exceptional candidates in all areas are encouraged to apply.

**Duties:** Teach and develop courses in electrical or computer engineering, develop an extramurally funded research program, publish outstanding work in leading scholarly journals, supervise graduate students, and provide department, college, and university service.

**Minimum qualifications:** An earned Ph.D. in Electrical or Computer Engineering, or a closely related discipline, with a strong research track record. All-But-Dissertation cases will be considered but degree must be earned before the date of hire. For complete duties and qualifications, and application instructions, refer to

**https://www.governmentjobs.com/careers/hawaiiedu and search for Position 84464.**

**Continuous recruitment:** Application reviews will begin on **July 15, 2022**, and will continue until the position is filled. The University of Hawaiʻi is an equal opportunity/affirmative action institution and encourages applications from women and minority candidates.

LUND
UNIVERSITY

**Professor in Analogue IC design**

Lund University is hiring a professor that is specialized in mm-wave and THz frequencies for future 5G-6G applications. The ideal candidate has outstanding research achievements in integrated wireless radio transmitters and receivers for RF / mm waves / THz communication, an interest in working with new technologies (including III-V technologies) in addition to silicon-based technologies. A high level of teaching skills, including a very good ability to conduct, develop and lead teaching and other educational activities on different levels and using a variety of teaching methods.


Connect with thousands of your colleagues in the **IEEE AuthorLab** today!

IEEE Collabratec™
Bright Minds. Bright Ideas.

collabratec.ieee.org

BY ALLISON MARSH

# Past Forward

## Pictures of a Planet

**Fifty years ago,** there was no good way to study the Earth's landmasses from space. Then, on 23 July 1972, NASA launched the Earth Resources Technology Satellite. ERTS carried a television-style camera system as well as an experimental Multispectral Scanner System (MSS), which capitalized on fiber optics to collect data at four different spectral bands—green, red, and two different bands of near-infrared. Everyone expected the TV camera to be the workhorse of the mission, but MSS turned out to be the winning technology.

Two days after launch, the Goddard Space Flight Center received its first MSS image: a view of Dallas. (In the false-color image, shown here, reds are vegetation and grays and whites are urban or rocky land.) The satellite orbited the Earth about 14 times a day, imaging the globe every 18 days. It gave scientists, land managers, and policymakers an unprecedented view of their planet. Over its five-and-a-half-year life span, the satellite provided more than 300,000 MSS images. In 1975, NASA renamed the ERTS program Landsat, and it is still going strong. The latest satellite, Landsat 9, launched on 27 September of last year, and the Landsat Next mission is already in the planning stages. ■

**FOR MORE ON THE HISTORY OF LANDSAT, SEE** spectrum.ieee.org/pastforward-jul2022

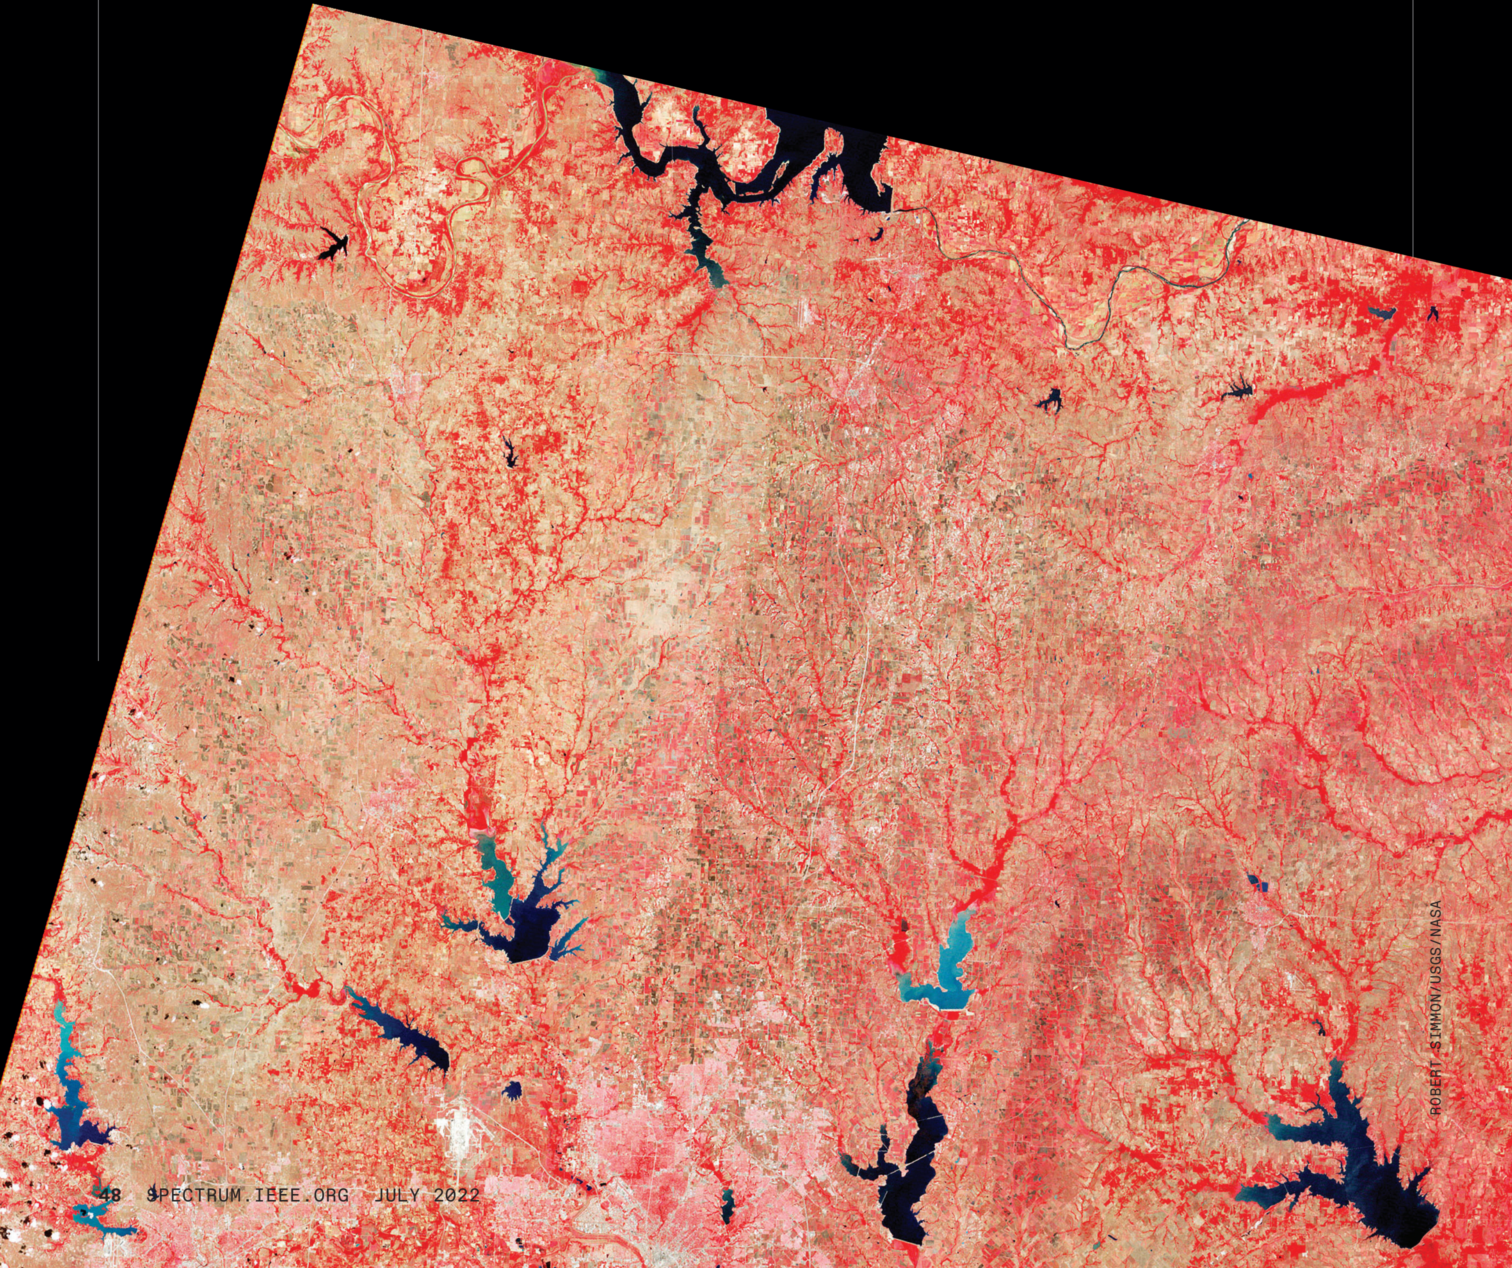

ROBERT SIMMON/USGS/NASA

BILLS
DON'T
STOP FOR
DISABILITY.
WHY
SHOULD
YOUR
INCOME?
Help protect your
livelihood and lifestyle.

MATLAB
HOME
PLOTS
APPS
EDITOR
New Script
New
Open
Find Files
Compare
Import Data
Save Workspace
New Variable
Open Variable
Clear Worksp
FILE
VARIABLE
C: ▸ Users ▸ Ainge ▸ AppData ▸ Local ▸ Mic
Editor - C:\Users\Ainge\AppData\Local\Microsoft\Windows\INetCac
Ready