News

COMPUTING

Open Source Has a "Maker-Taker" Problem

› Here are three ways to fix it

BY RINA DIANE CABALLAR

WordPress, the open-source Web content-management system that powers about 40 percent of the world's websites, is currently embroiled in a legal battle with WP Engine, a hosting provider for WordPress–built websites. Matt Mullenweg, a cofounder of WordPress, has accused WP Engine of violating WordPress trademarks. WP Engine hit back with a lawsuit.

The dustup is a high-profile example of what Dries Buytaert, founder of the open-source content-management system Drupal, calls the "maker-taker problem" in open-source software. In a blog post, he writes that "creators of open-source software ('Makers') see their work being used by others, often service providers, who profit from it without contributing back in a meaningful or fair way ('Takers')."

Chad Whitacre, head of open source at the application-monitoring-software company Sentry, says that "the point of the open-source licenses we have is this permissionless sharing—that's the blessing and the curse of open source. Everybody gets to use it, but it puts us precisely in this position of not being able to support that directly through the market."

The result, experts observe, is burnout for maintainers of open-source software and a crisis in its sustainability. Whitacre suggests there are three levers that can be pulled to solve the dual crises of open-source sustainability and maintainer burnout.

The crises are a symptom of how successful and important open-source proj-

Illustration by Dan Page

ects have become. Small, fledgling projects by enthusiastic and collaborative-minded software developers have grown into essential components of huge computer and Internet infrastructure. The Linux Foundation estimates that about 70 to 90 percent of today's apps use open-source software.

This increased reliance on open source is putting pressure on project maintainers, especially when it comes to providing prompt software updates and critical security fixes. In late 2021, a vulnerability in the widely used Log4j tool was one of the highest-profile security incidents in the open-source world. In 2023, a backdoor hack on a Linux compression tool was uncovered—a result of an attacker posing as a contributor to gain the maintainer's trust over two years, allowing the attacker to insert dangerous code into the tool. Both security flaws were in open-source projects maintained by small teams, yet the projects underlie massive enterprises.

Maintainers of open-source projects are also experiencing rising rates of burnout. According to a 2024 survey by Tidelift, a company that partners with and pays open-source maintainers to implement secure software-development practices, the top three things that respondents dislike about being maintainers are not being sufficiently compensated (or at all) for their work, feeling underappreciated, and that the work adds to their personal stress. It comes as no surprise, then, that more than half of maintainers have quit or have considered quitting.

So how can companies benefiting from open-source projects better support the community? Whitacre's three levers rely on commercialization, taxation, and social validation.

Commercialization is the traditional route, which involves finding business models to subsidize open-source projects. "You're not directly paying for open-source software, but you're paying for something else that is supporting or subsidizing the open-source work," Whitacre says. A classic example is making the software itself free but charging for support and services. "The key to the commercialization lever is you need something scarce to have a business. Open source is post-scarcity by definition [and] by intention, so you have to find something else scarce," he says.

Taxation is another lever to economically sustain open source. In Germany, for example, the Federal Ministry for Economic Affairs and Climate Action is financing the Sovereign Tech Fund to invest in "projects that benefit and strengthen the open-source ecosystem." Currently funded projects include FreeBSD, a Unix-like operating system; JUnit, a testing framework for the Java ecosystem; the PHP Foundation, which is behind the PHP programming language; and the OpenJS Foundation, which hosts a range of JavaScript projects.

Social validation is the final lever, according to Whitacre. Buytaert follows this approach with Drupal by implementing a credit system that recognizes and provides incentives to contributors for their efforts. Individuals and organizations who contribute to Drupal—whether that's through code, documentation, submitting case studies that demonstrate success with the software, or financial support—earn credits for visibility and advertising on Drupal's website, as well as early access, discounts, and sponsorships to events, among other benefits.

Whitacre himself is pulling the social-validation lever as a leader of the Open Source Pledge, a group working to directly pay maintainers. By joining the Pledge, companies pay an annual minimum of US $2,000 per full-time developer on staff to open-source maintainers of their choosing.

Members of the Open Source Pledge are also required to publish a blog post detailing their payments to maintainers, for accountability and awareness.

Since it began in October 2024, the Open Source Pledge has gained two dozen or so members, but those members are mostly "smaller, developer-focused startups highly aligned with open source to begin with," says Whitacre. The aim is to expand to larger enterprises.

Looking to the future of open-source software, Whitacre hopes to see more approaches similar to what Drupal is doing, as well as a commitment from both government bodies and the tech sector to steward open-source projects.

"With the Pledge, we're trying to get the money flowing right, but that's only half the equation," says Whitacre. "It's a very important half, and it's the part we need to start with and focus on right now. The other side of that is, how do we make sure that money actually has the impact we want? We need to unlock this same kind of commitment from the whole industry, broaden it, and get other companies to join us." ■

"The point of the open-source licenses we have is this permissionless sharing—that's the blessing and the curse of open source."

—CHAD WHITACRE, SENTRY

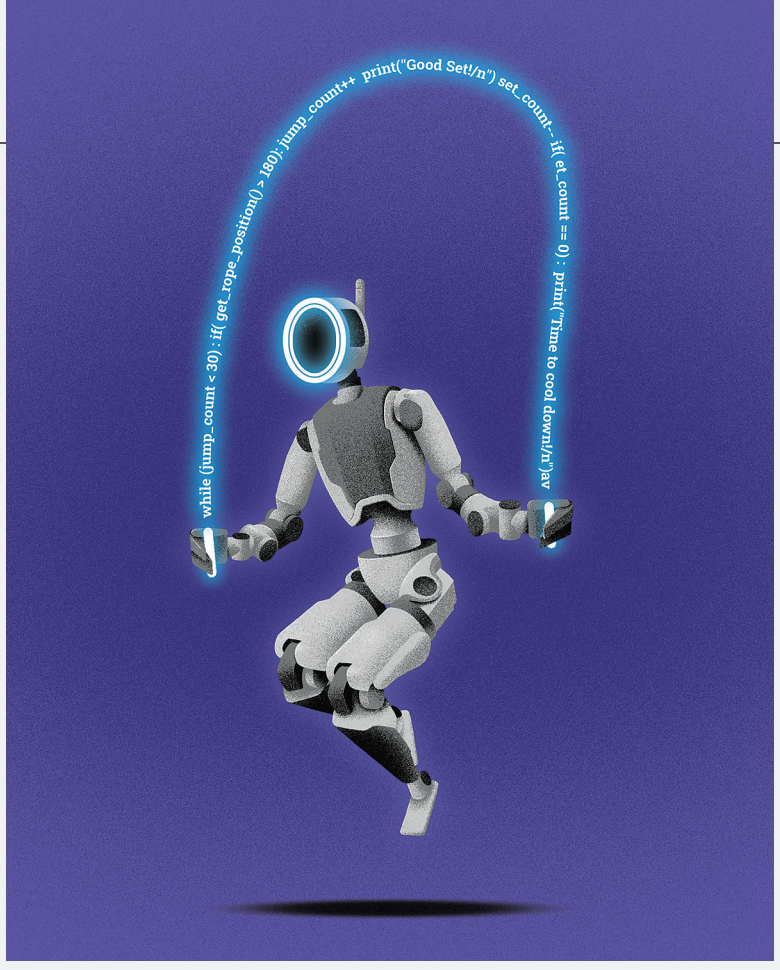

ROBOTICS

Boston Dynamics and Toyota Research Team Up on Robots

› The partnership will bring Atlas useful new skills

BY EVAN ACKERMAN

Boston Dynamics and the Toyota Research Institute (TRI) recently announced a new partnership to get closer to those perpetually overpromised general-purpose humanoid robots.

With its electric Atlas robot, Boston Dynamics already has a humanoid platform capable of advanced whole-body motions, but so far we've only seen Atlas do a few targeted manipulation tasks.

Meanwhile, TRI has been working on AI-based learning techniques to tackle a variety of complicated manipulation challenges, but they've been mostly stuck using stationary robotic arms. The partnership brings together AI teams from both organizations; Boston Dynamics gets some useful new skills for Atlas, and TRI gets a new hardware platform for its experiments. The rest of us get a hint of what a general-purpose humanoid robot might actually look like in the near future.

"We're seeing the limits of what you can do in tabletop manipulation, and we want to explore beyond that," says Russ Tedrake, vice president of robotics research at TRI. "In my mind, robot learning has opened up a brand new science—there's a new set of basic questions that need answering."

Progress in robotics used to be relatively small scale, but Tedrake believes the focus on getting humanoid robots to commercial viability has changed that. "Robotics has come into this era of big science, where it takes a big team and a big budget and strong collaborators," he says.

Scott Kuindersma, senior director of robotics research at Boston Dynamics, explains that Atlas is uniquely positioned to explore everything that humanoid robots can do, thanks to exceptionally flexible hardware that allows the robot to move in ways that humans can't. Atlas, for example, can twist its torso (and its head) all the way around.

TRI has been developing large behavior models (LBMs), which are analogous to large language models but were created for robots doing dexterous tasks in the physical world. Using only a few dozen human demonstrations combined with a natural-language description of the goal, a robot can learn and robustly execute a complex manipulation skill within a matter of hours. TRI's approach is based on diffusion, which is a generative AI technique that plans robot behaviors in an adaptive way that's particularly well-suited to human demonstrations.

For robots to learn from humans, they'll need to handle the fact that different humans may do the same task in different ways, and also that the same task may be done differently in different settings. LBMs should give general-purpose humanoid robots a foundation from which to take skills they learn from specific demonstrations and generalize those skills to make them real-world useful.

Or at least that's the goal, but TRI and Boston Dynamics are still puzzling through how to make it happen, says TRI's Tedrake. "Nobody has the beginnings of an idea of what the right training mixture is for humanoids. We have to figure it out."

Both TRI and Boston Dynamics stress that the collaboration's purpose is fundamental research rather than any commercial goals. But Kuindersma acknowledges that there will certainly be significant long-term commercial value. "It's not lost on either of our organizations that this is maybe one of the most exciting points in the history of robotics," says Kuindersma. "But there's still a tremendous amount of work to do." ■

Illustration by Luca D'Urbino

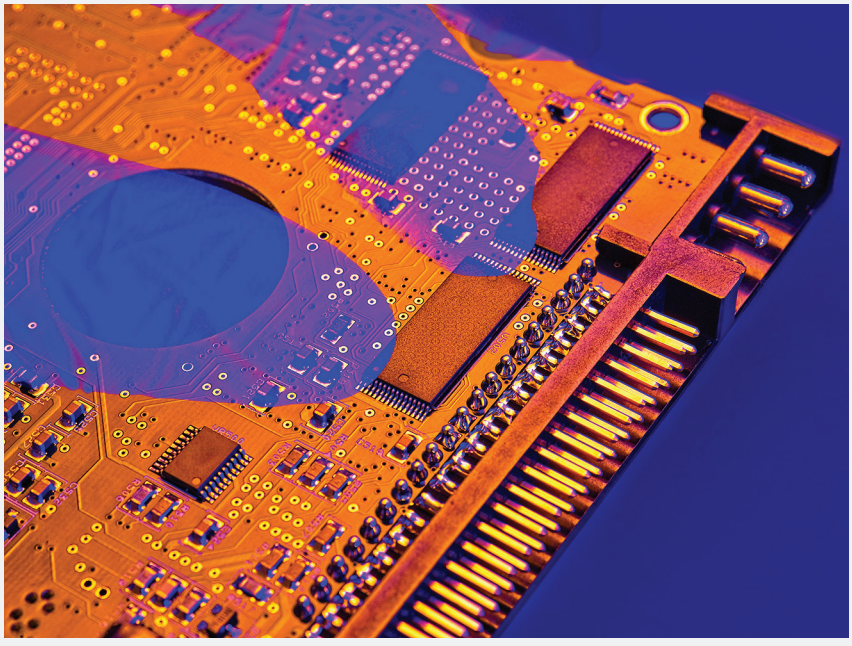

ARTIFICIAL INTELLIGENCE

Analog Startup Puts Generative AI on a Power Diet › Sagence AI promises large language models at 10 percent power

BY SAMUEL K. MOORE

Machine learning chips that use analog circuits instead of digital ones have long promised huge energy savings. But in practice they've mostly delivered modest savings, and only for modest-size neural networks.

Silicon Valley startup Sagence AI says it has the technology to bring the promised power savings to tasks suited for massive generative AI models. The startup claims that its systems will be able to run Meta's large language model Llama 2 70B on one-tenth the power of a Nvidia H100 GPU-based system, at one-twentieth the cost and in one-twentieth the space.

In simulations, a system made up of Delphis ran Llama 2 70B at

666,000

tokens per second while consuming 59 kilowatts.

"My vision was to create a technology that was very differentiated from what was being done for AI," says Sagence AI CEO and founder Vishal Sarin. When he founded the company in 2018, he recognized that power consumption would block mass adoption of AI—especially as models ballooned in size.

The core power-savings prowess for analog AI comes from two fundamental advantages: It uses some basic physics to do machine learning's most important math, and it doesn't have to move data around.

That math function is multiplying vectors and then adding up the result, called a multiply-accumulate operation. A decade ago, engineers realized that two foundational rules of electrical engineering did the same thing, more or less instantly. Ohm's Law—voltage multiplied by conductance equals current—does the multiplication if you use the neural network's "weight" parameters as the conductances. Kirchhoff's Current Law—the sum of the currents entering and exiting a point is zero—means you can add up all those multiplications just by connecting them to the same wire.

And in analog AI, the neural network parameters don't need to be moved from memory to the computing circuits—usually a bigger energy cost than computing itself—because they are already embedded within the computing circuits.

Sagence AI uses flash memory cells as the conductance values. The kind of flash cell typically used in data storage is a single transistor that can hold 3 or 4 bits, but Sagence AI has developed algorithms that let cells embedded in their chips hold 8 bits, which is the key level of precision for LLMs and other so-called transformer models. Storing an 8-bit number in a single transistor saves cost, space, and energy, says Sarin. Digital-memory technologies require multiple transistors per bit, so storing an 8-bit number would require 48 transistors.

Adding to the power savings is that the flash cells are operated in a state called "deep subthreshold," where they produce very little current. That wouldn't do in a digital circuit, because it would slow computation to a crawl. But because the analog computation is done

PETER BLOTTMAN PHOTOGRAPHY/ISTOCK

all at once, it doesn't hinder the speed.

Back in 2017 and 2018, a trio of startups went after a version of flash-based analog AI. Syntiant eventually abandoned the analog approach for a digital scheme that has put six chips in mass production so far. Anaflash and Mythic struggled, but stuck with it. Others, particularly IBM Research, have developed chips that rely on nonvolatile memories other than flash.

Analog AI has labored to meet its potential, particularly when scaled up to useful sizes. Flash cells have natural variations in their conductances, which might mean that the same number stored in two different cells will result in two different conductances. Worse still, these conductances can drift over time and shift with temperature. This noise can drown out the signal representing the result, and the noise can be compounded through the many layers of a deep neural network.

Sagence AI's solution, Sarin explains, is a set of reference cells on the chip and a proprietary algorithm that uses them to calibrate the other cells and track temperature-related changes.

Analog AI also requires digitizing the result of the multiply-accumulate operation in order to deliver it to the next layer of the neural network, where it must then be turned back into an analog voltage signal. These steps require analog-to-digital and digital-to-analog converters, which take up chip area and soak up power. According to Sarin, Sagence AI has developed low-power versions of both circuits.

Sagence AI's first product, to launch this year, will be geared toward vision systems, which are a lighter lift than server-based language models. "That is a leapfrog product for us, to be followed very quickly [by] generative AI," says Sarin.

The generative AI product would be scaled up from the vision chip mainly by vertically stacking analog AI chiplets atop a communications die. These stacks would be linked to a CPU die and to high-bandwidth memory DRAM in a single package called Delphi.

In simulations, Sagence AI claims, a system made up of Delphis ran Llama 2 70B at 666,000 tokens per second while consuming 59 kilowatts, versus 624 kW for an Nvidia H100–based system. ■

JOURNAL WATCH

Bluetooth Microscope Peers Inside Live Mice

Researchers have recently designed a prototype microscope that captures real-time imaging from inside live mice and transmits it via Bluetooth to a nearby phone or laptop.

Existing microscopes designed for internal study of living animals often require a wire attached to a monitoring device—and the animal—to transmit the data, which restricts the animal's natural behavior. Although the Bluetooth microscope, called the BLEscope—"BLE" stands for "Bluetooth low energy"—is currently 65 by 55 by 20 millimeters, the researchers are working to further miniaturize the components in order to study mice and other animals as they scurry about.

Arvind Pathak, a professor at the Johns Hopkins University School of Medicine, and his colleagues developed the BLEscope and described it on 24 September in *IEEE Transactions on Biomedical Engineering*.

Pathak says the BLEscope consumes less than 50 milliwatts of power and achieves a spatial resolution of between 5 to 10 micrometers. The microscope has two LED lights of different colors, green and blue, that create contrast during imaging.

In one experiment, Pathak and his colleagues added a fluorescent tracer into the blood of live mice and used the BLEscope to characterize blood flow within the animals' brains in real time. In another experiment, the researchers altered the oxygen and carbon dioxide ratios of the air inhaled by mice with brain tumors and observed blood-vessel changes in the tumors that were marked.

—Michelle Hampson

A Bluetooth-enabled microscope captures a fluorescent tracer in a mouse's tissue.

THE PATHAK LAB/JOHNS HOPKINS UNIVERSITY SCHOOL OF MEDICINE

ENERGY

Azerbaijan Plans Caspian–Black Sea Energy Corridor

› Link would move clean electricity from Caucasus to Europe

BY AMOS ZEEBERG

During COP29, the United Nations' climate-change conference, in November, the host country of Azerbaijan promoted a grand multinational plan to generate renewable electricity in the Caucasus region. The plan would send electricity nearly 1,200 kilometers west, under the Black Sea, and into energy-hungry Europe.

The transcontinental connection would first send up to 1.3 gigawatts of wind, solar, and hydropower generated across Azerbaijan and Georgia—as well as offshore wind power generated in the Caspian Sea—to the town of Anaklia, in Georgia, at the east end of the Black Sea. An undersea cable would then move the electricity across the Black Sea and deliver it to Constanta, in Romania, where it could be distributed further into Europe.

The scheme's proponents say this Caspian–Black Sea energy corridor will help decrease global carbon emissions, provide dependable power to Europe, modernize developing economies at Europe's periphery, and stabilize a region shaken by war. Organizers hope to build the undersea cable within the next six years at an estimated cost of US $3.1 to 3.6 billion.

To accomplish this goal, the governments of the involved countries must quickly circumvent a series of technical, financial, and political obstacles. "It's a huge project," says Zviad Gachechiladze, one of the corridor's architects and a director at Georgian State Electrosystem, the agency that operates the country's electrical grid. "To put it in operation [by 2030]—that's quite ambitious, even optimistic," he says.

The technical linchpin of the plan is the successful construction of the high-voltage direct-current (HVDC) submarine cable in the Black Sea. Laying the cable is a formidable task, considering that it would stretch across nearly 1,100 kilometers of water, most of which is over 2 km deep, and—since Russia's invasion of Ukraine—littered with floating mines. The longest existing submarine power cable, the North Sea Link, carries 1.4 GW across 720 km between England and Norway, at depths of up to 700 meters.

Building such connectors involves laying and stitching together lengths of heavy submarine power cable using specialized ships—the expertise for which lies with just two companies in the world: France-based Nexans and Italy-based Prysmian.

In an assessment of the Black Sea project's feasibility, the Milan-based consulting and engineering firm CESI determined that the undersea cable could be installed, and estimated that it could indeed carry up to 1.3 GW—enough to supply over 2 million European households.

To fill that pipe, countries in the Caucasus region would have to generate much more green electricity. Georgia would mostly use hydropower, which already generates more than 80 percent of the nation's electricity. "We are a hydro country. We have a lot of untapped hydro potential," says Gachechiladze.

But generating hydropower can also generate opposition because dams alter rivers and landscapes. Local opposition is "definitely a problem, and it has not been totally solved," says Salome Janelidze, a board member at the Energy Training Center, a Georgian government agency that promotes and educates around the country's energy sector. But "to me it seems it is doable," she says. "You can procure and construct if you work closely with the local population and see them as allies rather than adversaries."

For Azerbaijan, most of the electricity would be generated by wind and solar farms funded by foreign investment. Masdar, the renewable-energy developer of the United Arab Emirates government,

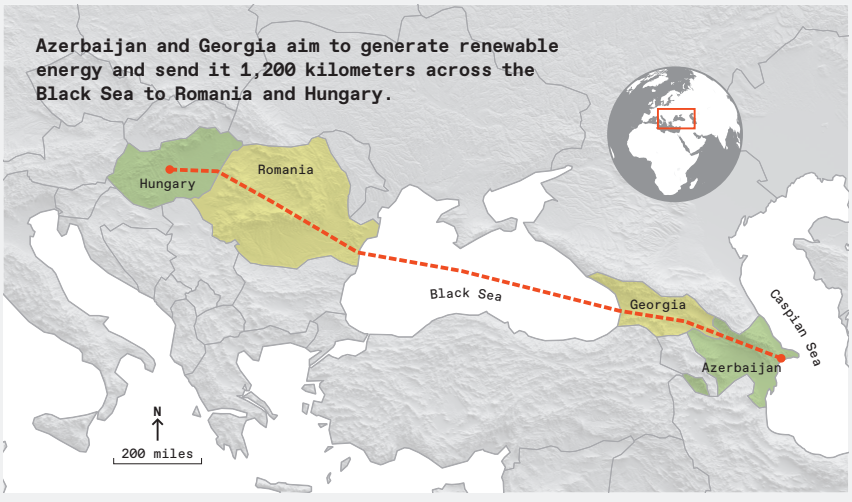

Map by Chris Philpot

has been investing heavily in wind power in the country. In June, the company broke ground on a trio of wind and solar projects in Azerbaijan with 1-GW capacity, and it intends to develop up to 9 GW more by 2030. Acwa Power, a Saudi Arabian power-generation company, plans to complete a 240-MW wind farm in the Absheron and Khizi districts of Azerbaijan this year. The company has also struck a deal with the Azerbaijani ministry of energy to install up to 2.5 GW of offshore and onshore wind.

CESI is currently running a second study to gauge the practicality of an even larger proposed energy corridor from the Caspian Sea to Europe, with a total transmission capacity of 4 to 6 GW. But that beefier interconnection will likely remain out of reach in the near term. "By 2030, we can't claim our region will provide 4 GW or 6 GW," says Gachechiladze.

In September, Azerbaijan, Georgia, Romania, and Hungary created a joint venture to shepherd the project. Those four countries had previously inked a memorandum of understanding with the European Union to develop the energy corridor.

They are now asking the E.U. to designate the cable a "project of mutual interest," which would make it an infrastructure priority for connecting the union with its neighbors. If selected, Gachechiladze says, the project could qualify for 50 percent grant financing. "It's a huge budget," he says. "It will improve drastically the financial condition of the project." An E.U. commissioner projected that the union might contribute €2.3 billion ($2.4 billion) toward building the cable.

There are security issues looming over the project. Shipping routes in the Black Sea have become less dependable since Russia's invasion of Ukraine. To the south, tensions between Armenia and Azerbaijan remain after a recent war and ethnic violence.

In order to improve relations between countries, many advocates of the energy corridor would like to include Armenia. "The cable project is in the interests of Georgia, it's in the interests of Armenia, it's in the interests of Azerbaijan," says Agha Bayramov, an energy geopolitics researcher at the University of Groningen, in the Netherlands. "It might increase the chance of them living peacefully together. Maybe they'll say, 'We're responsible for European energy. Let's put our egos aside.'" ■

The Falcon Neuro optical sensor [red box] spent two years on the International Space Station watching for lightning.

ASIM/EUROPEAN SPACE AGENCY

AEROSPACE

Neuromorphic Imager Maps Lightning Strikes From Space › Its improved resolution could lead to better storm forecasting

BY RACHEL BERKOWITZ

Our warming planet is beginning to feel the effects of a future with more severe thunderstorms and increased lightning strikes. But there's a lot that atmospheric scientists don't understand about how lightning works. Better lightning data could improve severe-weather forecasts and warnings, as well as help predict wildfires and other consequences.

A unique optical sensor called the Falcon Neuro, which just spent two years on the International Space Station, could help fill those gaps. Researchers at Western Sydney University, supported by the U.S. Air Force Research Lab, demonstrated the use of an event-based vision sensor to record lightning-strike details from above. The technology "can see things that a normal camera can't," says Gregory Cohen, the deputy director of Western Sydney University's International Centre for Neuromorphic Systems. The sensor offers higher resolution, lower cost, and lower data rates than other technologies.

Researchers don't exactly know how much energy is emitted at each point during the discharge process. Nor do they fully understand how lightning is initiated within thunderclouds, how it travels through a cloud, or how often it may happen. Better understanding and forecasting requires capturing features in cloud-to-cloud and cloud-to-ground lightning that are difficult to pin down.

Vanna Chmielewski, a research physical scientist at the U.S. National Oceanic and Atmospheric Administration's National Severe Storms Laboratory, is one scientist who would appreciate more observations. "We have global observations from the ground of where lightning strikes or transfers current, and we have a broad view from space," she says. But it's a challenge trying to align ground and space observations and understand what different instruments are seeing, Chmielewski says.

Lightning researcher Steven Goodman, emeritus senior scientist on the Geostationary Operational Environmental Satellites program, says that "by 2040 we'll have a ring of geostationary satellites mapping lightning around the planet." But returning reams of data from geostationary orbit is difficult, slow, and expensive.

That's where the Falcon Neuro instrument comes in. It's a pair of neuromorphic vision sensors, or event-based sensors. Instead of capturing an image frame by frame, these devices capture changes in light level within a scene. The raw output is a stream of "events" for each pixel—every instant at which the brightness change exceeds a certain threshold, with microsecond timing. The result is high-contrast, high-speed observations. "We get all the benefits that we need for space—low power, high speed, low data rate, high dynamic range," says Cohen.

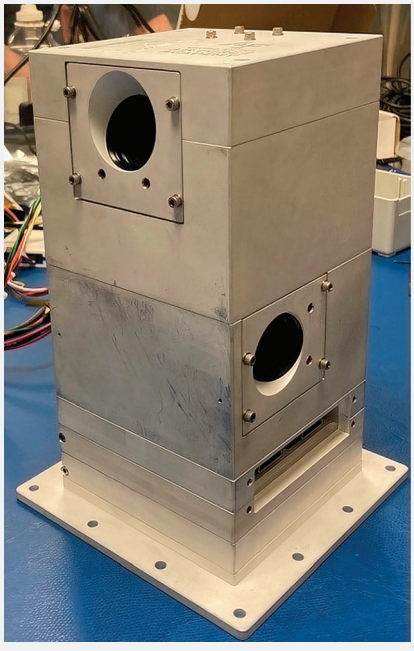

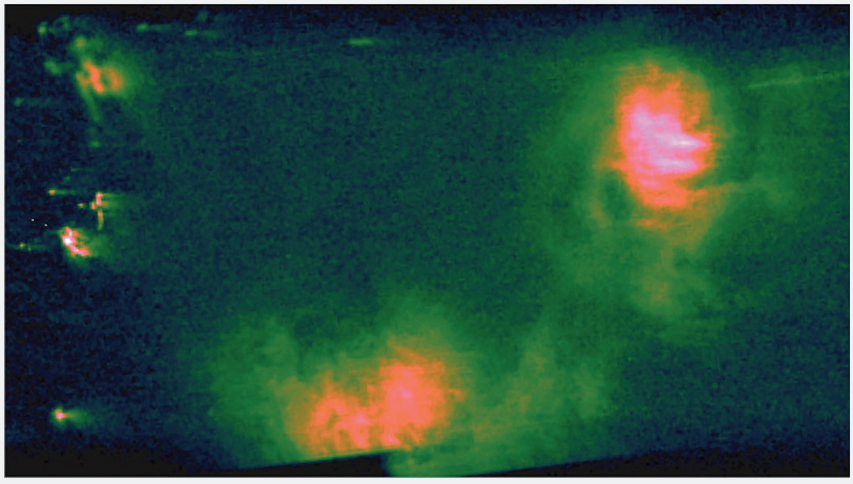

The Falcon Neuro sensor is able to spot lightning flashes [above] in quick succession, such as this pair that are within 5 seconds of each other. The instrument includes two cameras [bottom left]: Nadir [at top] and Ram, which look forward and down, respectively.

"By 2040 we'll have a ring of geostationary satellites mapping lightning around the planet."

—STEVEN GOODMAN, EMERITUS SENIOR SCIENTIST, GEOSTATIONARY OPERATIONAL ENVIRONMENTAL SATELLITES PROGRAM

Cohen, Geoff McHarg at the University of Iowa, and their colleagues designed the imager using two heavily modified sensors, one looking down and one looking forward. The imager gathers light across the visible and near-infrared spectrum from its perch aboard the ISS. The team developed a feature-finding algorithm that identifies probable lightning events in the raw data that's been processed into images with meter-scale resolution.

To confirm the presence of lightning in data from January 2022 to August 2023, the researchers compared suspected strikes to those detected by the ground-based radio-frequency sensors of the Global Lightning Detection Network (GLD360). They found that every time the GLD360 recorded a lightning flash, Falcon Neuro recorded a significant increase in clusters of events happening at the same time and place.

Moreover, the instrument routinely detected multiple lightning flashes in the same cloud, while the GLD360 recorded only a single flash—indicating that Falcon Neuro was capturing detailed features of lightning progression that were out of view of its ground-based counterpart. And these high-speed recordings—equivalent to 500 to 1,000 frames per second—generated only 3 to 4 megabits per second of data. (By comparison, a typical Starlink satellite user experiences data rates between 25 and 220 Mb/s).

Cohen adds that the demo uses camera components from the 1990s. He's working on future versions that can make high-speed lightning recordings at more wavelengths. ■

TOP: M. MCHARG, I. JONES ET AL./ACS; BOTTOM: U.S. AIR FORCE ACADEMY SPARC