IEEE
SPECTRUM

FOR THE TECHNOLOGY INSIDER | 09.19

A human heart beats 42 million times a year. Making a machine that can pump so reliably has been a challenge

The Magnetic Heart

Now, a group led by Daniel Timms is about to start human trials on the most impressive artificial heart since 2006
P. 22

IEEE

CONTENTS_09.19

IEEE SPECTRUM

THE MAGLEV HEART

A machine with one moving part could give heart-failure patients new life.

BY NICHOLAS GREATREX, MATTHIAS KLEINHEYER, FRANK NESTLER & DANIEL TIMMS *Page 22*

On the cover
Photograph for *IEEE Spectrum* by Peter Adams

IEEE SPECTRUM
(ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2019 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 56, Issue No. 9. 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, c/o 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 http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

ILLUSTRATION BY Bryan Christie Design

DANIEL TIMMS

TAKING LESSONS TO HEART

When Daniel Timms was growing up in Brisbane, Australia, he spent many hours helping his father build wild contraptions featuring pumps and waterfalls [above]. His father, a plumber with a passion for invention, taught Timms about fluid dynamics and also instilled in him "a practical attitude toward getting things done," Timms says.

In 2001 Timms's father was diagnosed with a condition that would gradually rob his heart of its ability to pump blood throughout his body. That's when Timms, who was then getting his Ph.D. in biomedical engineering, began working on a revolutionary design for an artificial heart [see "The Maglev Heart," p. 22]. He enlisted his father in the effort, and the two began tinkering with prototypes in the backyard shed. With pipes and valves from a local hardware store, they built a rudimentary model of the human cardiovascular system so that they could hook up prototype heart pumps for testing.

After Timms got his Ph.D., he went to Brisbane's Prince Charles Hospital and convinced physicians in the cardiology unit to clear out a room, which he turned into an engineering lab. Just down the hall was the intensive care unit (ICU), where his father regularly ended up as his heart problems worsened. Timms would drop his tools and walk down the hall to visit.

The last time his father was admitted to the ICU, in 2006, it was Timms who drove him to the hospital. Timms was supposed to get on a plane to Germany the following day, where he was to meet with potential collaborators. Timms asked his father whether he should cancel the trip. "He said, 'You've got to get there; this is what we've been working for,'" Timms remembers. His father passed away a few days later. But the trip did lead to a fruitful collaboration, which led to other partnerships in Japan, Taiwan, and the United States. Today Timms's company, Bivacor, has its headquarters in Houston, where it's preparing for clinical trials of its artificial heart. Timms is sure his father would be pleased with the outcome of their backyard tinkering. ■

CONTRIBUTORS_

Swarup Bhunia

Bhunia is a professor at the University of Florida. He and Soumyajit Mandal of Case Western Reserve University are developing portable instrumentation for testing medicines, as they describe in "Countering Counterfeit Drugs" [p. 38]. Bhunia's interest in the topic stems from personal experience. While visiting family in India, he says, "I had to buy medicines from street pharmacies. I did that a couple of times, but things were not working."

Zoë Christie

Christie is the founder and creative director of Bryan Christie Design, a 3D-illustration studio that specializes in science and technology. She and Joe Lertola, the studio's information graphics director, created this month's cover illustration of Bivacor's artificial heart. "Some things I've illustrated are very experimental and never amounted to anything," Christie says. "But the first 3D illustration I ever worked on was an artificial heart. That was 20 years ago, and the technology is still moving forward."

Adam Czajka

Czajka, a professor at the University of Notre Dame, has been mesmerized by the human iris for decades. In "The Eyes Have It" [p. 44], he, Mateusz Trokielewicz, and Piotr Maciejewicz explain that much of what's been assumed about the human iris is wrong, which has important implications for the security of iris-recognition systems. With a better understanding of how the iris changes during life (and beyond), the researchers are now working to make these systems more inclusive.

Allison Marsh

Marsh is codirector of the Ann Johnson Institute for Science, Technology & Society at the University of South Carolina. She writes *IEEE Spectrum*'s Past Forward column, which tells the story of technology through historical artifacts. In this issue, she considers the electric hot comb [p. 51], a popular gadget for straightening and shaping hair. It's a history she can relate to. In her big-hair days of high school, Marsh carried a butane-powered curling iron to restyle her bangs after gym class.

John Voelcker

Voelcker worked at *Spectrum* in the late 1980s, at *Green Car Reports* in the 2010s, and at a dozen media and tech companies in between. In his spare time, he tinkers with old British cars, a retro pursuit he balances with a focus on the future of transportation. The latter interest recently took him to Porsche's R&D center in Weissach, Germany, where he witnessed *really* fast electric-vehicle charging [see "Porsche's Fast-Charge Power Play," p. 30].

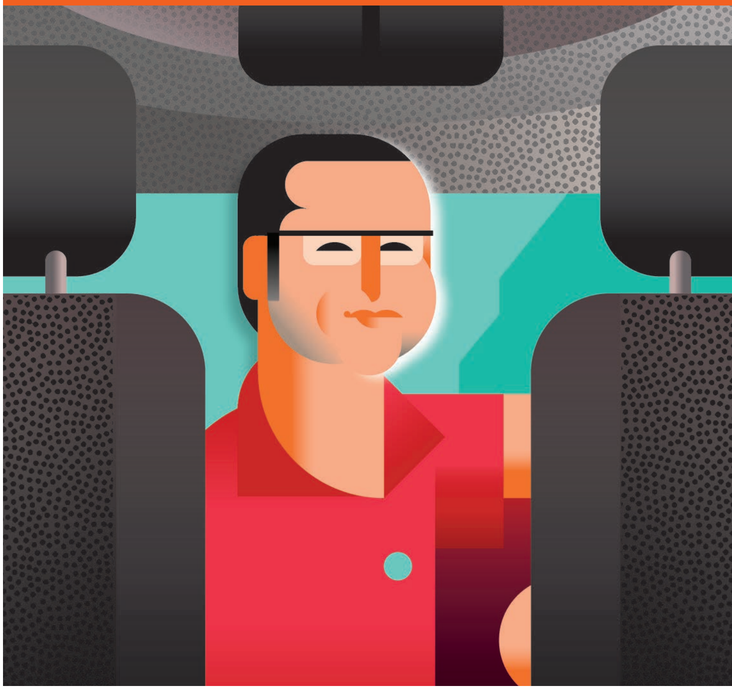

AI Engineers: The Autonomous-Vehicle Industry Wants You

Cruise's AI chief, Hussein Mehanna, talks jobs, careers, and self-driving cars

Three years ago, Cruise, an autonomous-vehicle startup acquired by General Motors, had about 50 employees. At the beginning of 2019, the head count at its San Francisco headquarters—mostly software engineers working on projects involving machine learning and artificial intelligence—hit around 1,000. Now that number is up to 1,500, and likely to reach about 2,000 by year-end, sprawling into a building that had housed Dropbox. And that's not counting the 200 or so tech workers that Cruise is aiming to install in a Seattle satellite development center and a handful of others in Phoenix and Pasadena, Calif. ¶ Cruise's recent hires aren't all engineers—it takes more than engineering talent to manage operations. And there are hundreds of so-called safety drivers that are required to sit in the 180 or so test vehicles whenever they roam San Francisco. But that's still a lot of AI experts to be hiring in a time of AI engineer shortages. ¶ Hussein Mehanna, head of artificial intelligence and machine learning at Cruise, says the company's hiring is on track, due to the appeal of the challenge of autonomous vehicles. Mehanna himself joined Cruise in May 2019 from Google, where he was director of engineering at Google Cloud AI. ¶ Mehanna has been immersed in AI and machine-learning research since his graduate studies in speech recognition and natural-language processing at the University of Cambridge, in the United Kingdom. I sat down with him to talk about his career, the challenges of recruiting AI experts, and autonomous-vehicle development in general. [*Editor's note: This interview has been condensed and edited for clarity.*]

Tekla S. Perry: When you were at Cambridge, did you think AI was going to take off ?

Hussein Mehanna: No. I do recall in 2003 that my supervisor and I were wondering if neural networks could help at all in speech recognition. Now neural networks have dominated vision, speech, and language processing. But that boom started in 2012.

I didn't expect it, but I certainly aimed for it when I was at Microsoft, where I deliberately pushed my career toward machine learning instead of big data, which was more popular at the time. And I aimed for it when I joined Facebook. In the early days, Facebook wasn't that open to Ph.D.s, or researchers. It actually had a negative sentiment about researchers. And then Facebook shifted to becoming one of the key places Ph.D. students wanted to join.

T.P.: Is it getting harder or easier to find AI engineers to hire, given the reported shortages?

H.M.: There's a mismatch between job openings and qualified engineers, though it is hard to quantify it.

Here at Cruise, demand for AI talent is just growing and growing. It might be saturating at other kinds of companies that are leveraging more traditional applications—ad prediction, recommendations.

The autonomous-vehicle problem is the engineering challenge of our generation. There's a lot of code to write, and if we think we are going to hire armies of people to write it line by line, it's not going to work.

Sometimes people worry that AI is taking jobs. It is taking some developer jobs, but it is generating other jobs as well, protecting developers from the mundane and helping them build software faster and faster.

T.P.: Where are you looking as you try to find a thousand or so engineers to hire this year?

H.M.: Because autonomous-vehicle technology is the new frontier for AI, the number of people with both AI and AV experience is quite limited. So we are acquiring AI experts wherever they are. You don't have to be an AV expert to flourish in this world.

There are endless applications to be developed over the next few decades. Even if we can get a car to drive safely, there's the question of, How can we tune the ride comfort and then apply it to different cities, different vehicles, different driving situations? I can see how I can spend a lifetime trying to solve this problem. —TEKLA S. PERRY

An extended version of this article appears on our View From the Valley blog.

POST YOUR COMMENTS at https://spectrum.ieee.org/aiengineers0919

CORRECTION: Due to an editing error, the August Hands On article, "Making Machine Learning Arduino Compatible," referenced "8-bit automatic-voice-recognition processors." This should have read "8-bit AVR processors."

ILLUSTRATION BY Maria Picassó i Piquer

NEWS

MARCELO MAGNASCO/ROCKEFELLER UNIVERSITY AND DIANA REISS/HUNTER COLLEGE

ELEPHANTS, DOLPHINS, AND CHIMPS NEED THE INTERNET, TOO

A new initiative promotes Internet communication among smart animals

People surf it. Spiders crawl it. Gophers navigate it. Now, a leading group of cognitive biologists and computer scientists want to make the tools of the Internet accessible to the rest of the animal kingdom.

Dubbed the Interspecies Internet, the project aims to provide intelligent animals such as elephants, dolphins, magpies, and great apes with a means to communicate with one another and with people online. And through artificial intelligence, virtual reality, and other digital technologies, researchers hope to crack the code of all the chirps, yips, growls, and whistles that underpin animal communication.

Oh, and musician Peter Gabriel is involved.

"We can use data analysis and technology tools to give nonhumans a lot more choice and control," the former Genesis front man said at the inaugural Interspecies Internet Workshop, held in July in Cambridge, Mass. "This will be integral to changing our relationship with the natural world."

Eighteen years ago, Gabriel visited a primate research center in Atlanta, where he jammed with two bonobos, a male named Kanzi and his half-sister Panbanisha. "It blew me away," he says.

WHACK-A-FISH: Foster, a dolphin at the National Aquarium in Baltimore, plays an interactive game. Optical sensors detect the dolphin's touch as fish move onscreen.

Gabriel eventually teamed up with Internet protocol co-inventor Vint Cerf, cognitive psychologist Diana Reiss, and Internet of Things pioneer Neil Gershenfeld to propose building an Interspecies Internet. For the past six years, the architects of this Dolittlesque initiative embarked on two small pilot projects, one for dolphins and one for chimpanzees.

At her Hunter College lab in New York City, Reiss developed what she calls the D-Pad—a touch pad for dolphins. Joining forces with Rockefeller University biophysicist Marcelo Magnasco, she invented an optical-detection system in which images and infrared sensors are projected from an underwater viewing gallery onto a glass panel, creating an aquatic touch screen through which dolphins can play specially designed apps, including one dubbed Whack-a-Fish.

Meanwhile, in the United Kingdom, Gabriel worked with Alison Cronin, director of the ape rescue center Monkey World, to test the feasibility of using FaceTime with chimpanzees.

The chimps engaged with the technology, Cronin reported at the workshop. However, our hominid cousins proved as adept at video-telephonic discourse as my 3-year-old son is at video chatting with his grandparents—which is to say there were a lot of silly games and not much meaningful conversation. The buggy, rudimentary attempt at interspecies online communication shows that building the Interspecies Internet will not be as simple as giving out Skype-enabled tablets to smart animals.

"There are all sorts of problems with creating a human-centered experience for another animal," says Gabriel Miller, director of research and development at the San Diego Zoo.

Miller has been working on animal-focused sensory tools such as an "Elephone" (for elephants) and the "Joy Branch" (for birds), but it's not easy to design efficient interactive systems for other creatures—and for the Interspecies Internet to be successful, Miller points out, "that will be super foundational."

Researchers are making progress on natural-language processing of animal tongues. Through a nonprofit organization called the Earth Species Project, former Firefox designer Aza Raskin and early Twitter engineer Britt Selvitelle are applying deep-learning algorithms developed for unsupervised machine translation of human languages to fashion a Rosetta Stone–like tool capable of interpreting the vocalizations of whales, primates, and other animals.

Inspired by the scientists who first documented the complex sonic arrangements of humpback whales—a discovery that ushered in the modern marine conservation movement—Selvitelle hopes that an AI-powered animal translator can have a similar impact on environmentalism today.

"A lot of shifts happen when someone who doesn't have a voice gains a voice," he says.

Verification and validation remain a challenge with this sort of AI software. Normally, machine-learning algorithms are benchmarked against a human expert, but who is to say if a cybernetic translation of a sperm whale's clicks is accurate or not?

According to primatologist Sue Savage-Rumbaugh, expertly trained bonobos could serve as bilingual interpreters, translating the argot of apes into the parlance of people, and vice versa.

Not just any trained ape will do, though. They have to grow up in a mixed *Pan/Homo* environment, as Kanzi and Panbanisha did. Unlike all other research primates, those apes thus grew up versed in both bonobo and human languages.

Panbanisha died in 2012, but Kanzi, nearly 39, is still going strong, and could in theory be recruited to check the accuracy of any Google Translate–like app for bonobo hoots, barks, grunts, and cries. And if wild bonobos in Central Africa can be coaxed to gather around a computer screen, Savage-Rumbaugh is confident Kanzi could communicate with them that way. "It can all be put together," she says. "We can have an Interspecies Internet."

That's music to the ears of Jeremy Coller, a private equity specialist whose foundation partially funded the Interspecies Internet Workshop. At the workshop, the foundation announced the creation of the Coller Doolittle Prize, a US $100,000 award to help fund further related research. A working group also formed to synthesize plans for the emerging field, facilitate future event planning, and guide testing of shared technology platforms.

Why would a multimillionaire with no background in digital communications or psychology want to back the initiative? For Coller, the motivation boils down to empathy: "If I can have a chat with a cow," he says, "maybe I can have more compassion for it." —ELIE DOLGIN

SOUND STICK: Samson, a hyacinth macaw, can play distinct sounds by perching on and manipulating this device, called the Joy Branch.

An extended version of this article appears in our Tech Talk blog.

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/animalinternet0919

RÉBECCA KLEINBERGER AND GABRIEL MILLER

AN AUTOMATED OFFSHORE FISH FARM COMES TO NORWAY

Giant remote-controlled pens will tend to millions of salmon

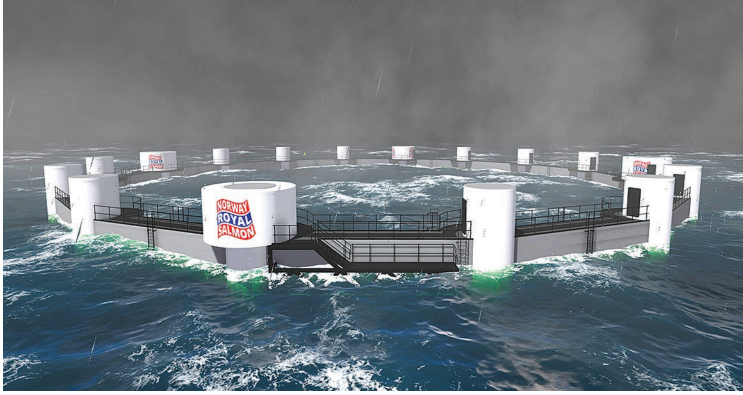

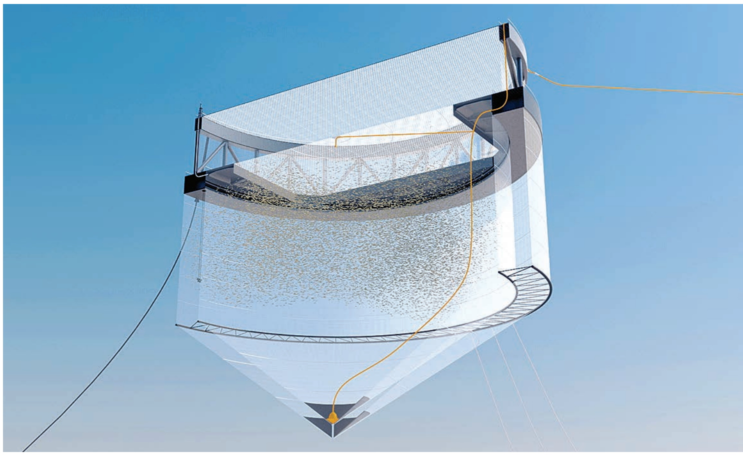

Tucked within Norway's fjord-riddled coast, nearly 3,500 fish pens corral upwards of 400 million salmon and trout. Norway exports more such fish than any other country in the world (1.1 million tons in 2018), and farmed salmon is one of its top three exports. With the global industry expected to quintuple by 2050, farmed salmon is a fine kettle of fish.

But raising salmon is not without challenges. Feeding them is expensive. Parasitic crustaceans called sea lice (*Lepeophtheirus salmonis*) attach to their bodies and graze on skin, blood, and mucus. If the lice don't kill the fish, some delousing methods, such as flushing the fish with water, might. About 15 percent of farmed salmon die annually, according to Norway Royal Salmon.

Now, a new remote-controlled fish pen—the first of its kind designed for the tempestuous waters of the open ocean—could help Norway meet the growing demand for salmon, reduce the cost of feed, and prevent deaths from sea lice.

In 2020, the first pen will be tested off the coast of the county of Troms, near a small island that will help abate the rough waves. If the field trials go well, the pens could be moved further out to sea.

Instrumented with wireless gauges, sensors, and cameras, the sea pen will allow workers onboard a nearby barge to monitor the fish and the system's automatic feeders. The feeding system reduces the energy cost of feeding by 50 percent compared with conventional methods. And because these open-ocean pens will keep salmon 10 to 40 meters beneath the surface, below the sunlit zone where sea lice and algae thrive, they could reduce or eliminate the need for delousing operations and significantly lower mortality rates.

The system, developed by Arctic Offshore Farming (owned by Norway Royal Salmon), resembles a giant fish basket. The cylindrical pen has a diameter of 79 meters and consists of a top and a bottom section. Each has a buoyant pontoon ring. When submerged, only the upper pontoon is visible on the ocean surface. Salmon live in the bottom section, in a large net that drapes to a depth of 40 meters from the lower pontoon. A net prevents the fish from swimming into the top section and can be removed to add fish or harvest them.

Fish need air to fill their swim bladders, so the pen has four decompressors that create air pockets beneath the water. Camera and oxygen sensors monitor the air pockets and automatically engage decompressors to keep them full at all times. The pontoons keep the pen stable, says Klaus Hatlebrekke, chief operating officer for market and business development at Norway Royal Salmon. "Even in a model test of a 50-year storm, we were not able to disrupt the air pockets," he says.

As with traditional salmon farming, the salmon will be raised from eggs to juveniles, called smolts, onshore in freshwater hatcheries. After a year, the smolts are transferred to a saltwater sea farm in a fjord. Under conventional methods, they would stay in the fjord and be raised until they weighed 5 kilograms. But with the new system, fish weighing 1.5 kg will be moved to the open-ocean pens, where they'll remain for another 10 to 11 months before they're harvested.

Using a local area network, the sensors transmit data to an onboard server that's connected via a fiber-optic cable to a crewed feed barge

TOP: NORWAY ROYAL SALMON; BOTTOM: ABB

SURF THE WAVES: A single pen built by Arctic Offshore Farming can hold 600,000 adult salmon. Fish are harvested from the bottom section of the two-tiered system.

stationed about 400 meters from the farm. One barge can monitor a cluster of fish farms and restock them with feed pellets every 7 to 14 days. With a volume of 120,000 cubic meters, each pen can hold up to 600,000 full-grown salmon at a time.

Standard feeding systems blow feed pellets through air hoses floating on the water's surface. But this approach doesn't work in the open ocean, where waves and winds would scatter the feed out to sea. Instead, the new feeding system automatically releases the feed underwater one to three times per day, allowing currents to distribute the pellets. Cameras allow crew members on the barge to see where the fish are located and release feed in those areas, reducing waste.

Although jobs in salmon fishing are about to get more technical, they don't have to be more complicated, says Lars Andersen, a sales specialist in aquaculture at ABB. The company built an interface for the salmon pen that displays the controls and safety systems in a simple dashboard.

As for the fish, their experience will more closely represent a life lived in the wild, says Hatlebrekke. Wild salmon that begin their journey in freshwater rivers spend only a few weeks in the sheltered fjord waters before they swim out to the open ocean. As farmed salmon operations incorporate more technology, the lives of farmed fish could mimic this path from freshwater to open sea, for a more natural existence.

—TRACY STAEDTER

A version of this article appears on our Tech Talk blog.

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/salmonfarm0919

HOW YOUTUBE LED TO GOOGLE'S CLOUD-GAMING SERVICE

The tech that made YouTube work everywhere promises to do the same for games

When Google's executives floated a vision for the Stadia cloud-gaming service, which could make graphically intensive gaming available on any device, they knew the company wouldn't have to build all the necessary technology from scratch. Instead, the tech giant planned to leverage its expertise in shaping Internet standards and installing infrastructure to support its YouTube video service for more than a billion people worldwide.

When Stadia debuts later this year, customers will be able to start gaming almost instantly by launching a simple client program that runs on Chromecasts, Chromebooks, PCs, and smartphones.

"Our vision is to have Stadia available on all devices that stream YouTube—a truly platform-agnostic service," says Majd Bakar, vice president of engineering for Stadia.

Cloud gaming has stricter technical requirements than streaming videos do. Video-streaming services such as YouTube and Netflix need to deliver video only when a person presses Play. For Stadia, a user's device must also perform additional processing to handle inputs from a player's controls.

Whereas live video streaming may have 500 milliseconds (half a second) to manage possible network glitches without noticeable interruptions, Stadia could have just 16 milliseconds or less to deliver a smooth experience for real-time interactive gaming.

Google harnessed technologies and infrastructure developed for YouTube—and engineered a few new ones—to prepare for Stadia's planned November 2019 launch in the United States, Canada, and a dozen European countries.

To use Stadia, Google recommends a minimum Internet download speed of 10 megabits per second to experience a game with 720p resolution at 60 frames per second. The service requires a minimum upload speed of 1.5 Mb/s.

Delivering high-intensity graphics over such limited bandwidth would not be possible without video codecs, which compress digital video files into smaller files. Compressed video files require less bandwidth and less time to transfer between a server and a person's device. That makes a huge difference for video streaming, which represents more than half of all downstream Internet traffic. YouTube alone makes up more than 10 percent.

Like YouTube, Stadia will lean heavily on video codecs to compress graphics into chunks of data that are easier to deliver. One called H.264 is a common compression standard for popular video-streaming services, Blu-ray discs, and HDTV broadcasts.

Google has developed a competitor called VP9 that delivers HD and 4K video streaming on YouTube with half the bandwidth of other codecs. The Stadia team has created "purpose-built custom hardware accelerators" that perform speedy encoding at scale based on both of those video codecs, Bakar says.

Another technology behind Stadia's cloud gaming is a variety of transport protocols—including basic protocols such as the Transmission Control Protocol (TCP)—that pass data between Internet-connected devices. One protocol used by Stadia, called WebRTC, arose from an open-source project supported by Google, Mozilla, and Opera. WebRTC allows software developers to build real-time video and audio communication into Web browsers such as Google Chrome and apps such as Google Hangouts and Duo.

Stadia will also rely on Google's QUIC (Quick UDP Internet Connections) protocol, which has reduced connection times and minimized delays in transmission in comparison with TCP. Google's Chrome browser and various apps already use QUIC for more efficient data transmission. But the protocol's capability to deliver data with less latency should also make a big difference for online gaming.

Another speed boost will come from Google's development of a congestion-control algorithm called BBR (Bottleneck Bandwidth and Round-trip propagation time). This algorithm can accurately measure Internet traffic and regulate how much data it puts into a network at any given time.

"Stadia's adaptive streaming technology adjusts to network quality in real time," Bakar says. "Alongside techniques like BBR, we are able to detect network impairments such as congestion prior to them happening."

These video codecs, Internet protocols, and congestion-control algorithms help Google make the best use of the existing bandwidth available in today's fiber-optic cables that connect data centers to customers. And like many tech giants, Google has invested in private Internet infrastructure so that it can more quickly deliver online services. In fact, Google has shared or complete ownership of approximately 8.5 percent of the world's submarine cables. And if there is a physical heart for Google's services, it's the company's 16 huge data centers, located in the Americas, Europe, and Asia.

Last but not least, Google has placed more than 7,500 edge nodes, which are Stadia servers installed in the networks of Internet service and network providers. Those edge nodes represent the Google infrastructure closest to customers.

Stadia will use a custom-built AMD GPU (graphics processing unit) that can deliver 10.7 teraflops of performance. That compares very favorably with the graphics cards in traditional video-game consoles such as the Xbox One X, with 6.0 teraflops, and the PlayStation 4 Pro's 4.2 teraflops—and Stadia could even leverage many graphics cards at once to supercharge gaming experiences.

Google's success in supplying YouTube and other Google services to billions of customers seems to have yielded useful lessons for the company's move into cloud gaming. "As long as these devices have good Internet connectivity and are able to decode high-quality video, they can handle Stadia," Bakar says. —JEREMY HSU

A version of this article appears in our Tech Talk blog.

REAL TIME: Google must reduce latency to 16 milliseconds for online gaming, compared with 500 ms of latency for streaming live video.

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/onlinegaming0919

GOOGLE

U.S. AIRLINE ORDERS FIRST PASSENGER ELECTRIC PLANE

The battery-powered nine-seater aircraft will enter service in 2022

Electric aviation took a big step forward in June when a Massachusetts-based airline announced it had placed the world's first order for a commercial all-electric passenger airplane. The Alice, a three-engine, battery-powered airplane that can fly up to 1,000 kilometers on a single charge, will be delivered to Cape Air in 2022.

The Alice, manufactured by the Israel-based startup Eviation Aircraft,

EVOLUTION OF FLIGHT: The Alice commuter aircraft is designed for flights of less than 1,000 kilometers. It has three propellers—one in back and two at the wingtips—and a 3,700-kilogram battery.

has not yet been certified by the U.S. Federal Aviation Administration. However, the company's e-airplane "could be certified right now to fly," insists Lior Zivan, Eviation's chief technology officer. The company is "anticipating full certification by 2022."

The Alice will be powered by a 900-kilowatt-hour lithium-ion battery manufactured by the South Korean company Kokam Battery. (For comparison, the Tesla Model 3 electric car uses a 50- to 75-kWh battery pack.)

Cape Air flies to popular vacation destinations, including Cape Cod and Martha's Vineyard, and serves a total of 35 cities with routes in the Caribbean and in the U.S. Northeast, Midwest, and Montana. According to Trish Lorino, Cape Air vice president of marketing and public relations, the company's historic order of Eviation's Alice aircraft "makes sense for us because we are a short-haul carrier." Lorino notes that "for 30 years, we have specialized in serving short-haul routes, particularly to niche and island destinations."

The carrier currently operates 88 Cessna 402s and four Islander planes (all of which seat nine passengers) made by the British company Britten-Norman. The nine-seater Alice e-aircraft thus fits within the Cape Air fleet's general size and passenger capacity.

Although the carrier has not yet decided which routes will feature the Alice, company officials anticipate keeping the plane close to the company's Massachusetts headquarters. "Short-haul routes in our 'backyard' such as Nantucket, Martha's Vineyard, and Provincetown would be the likely routes," Lorino says.

Eviation CEO Omar Bar-Ohay showcased the Alice at the Paris Air Show in June and spoke about the differences inherent in designing an all-electric airplane and a conventional, petroleum-fueled plane. For example, no fuel is burned during flight, so the plane's takeoff weight (6,350 kilograms or 14,000 pounds) is more or less its landing weight (and the battery accounts for 3,700 kg).

Each of the Alice's three motors has only one moving part, compared with about 10 in a standard (petroleum-fueled) reciprocating engine. "Obviously, electric propulsion has a major advantage in both reliability and maintenance," Zivan says.

The e-aircraft's three engines include two "pusher" motors mounted at the rear ends of the wingtips and another pusher motor mounted at the rear of the plane. All three are designed to have double or triple redundancy in their components. As for the electrical system, "the battery is designed in such a way that any malfunction or failure will result in a minimal reduction in the capacity, if any," Zivan says.

Because the Alice relies only on electric charge, the cost of operating the plane is expected to be lower than for its petroleum-fueled counterparts. And the noise emitted by a plane with no internal combustion engines is also lower. This is especially true for the Alice, given its ability (unique to e-aircraft) to vary its propeller speeds to compensate for crosswinds and to lower cabin noise.

As an early standard-bearer in electric passenger flight, Cape Air says its decision to purchase the Alice was also partly motivated by the company's "deep sense of social responsibility," Lorino says. The company's headquarters is 100 percent solar powered, she says, and the company hopes to use sustainable energy sources to charge its fleet of e-airplanes. The number of new electric aircraft that will join Cape Air's fleet has not been finalized.

"Our hope is that electric-powered flight is a reality in the next decade and that there is adoption from the public to view this as a viable, natural form of transportation," she says. —MARK ANDERSON

A version of this article appears in our Energywise blog.

↗ POST YOUR COMMENTS at https://spectrum.ieee.org/electricplane0919

EVIATION AIRCRAFT (2)

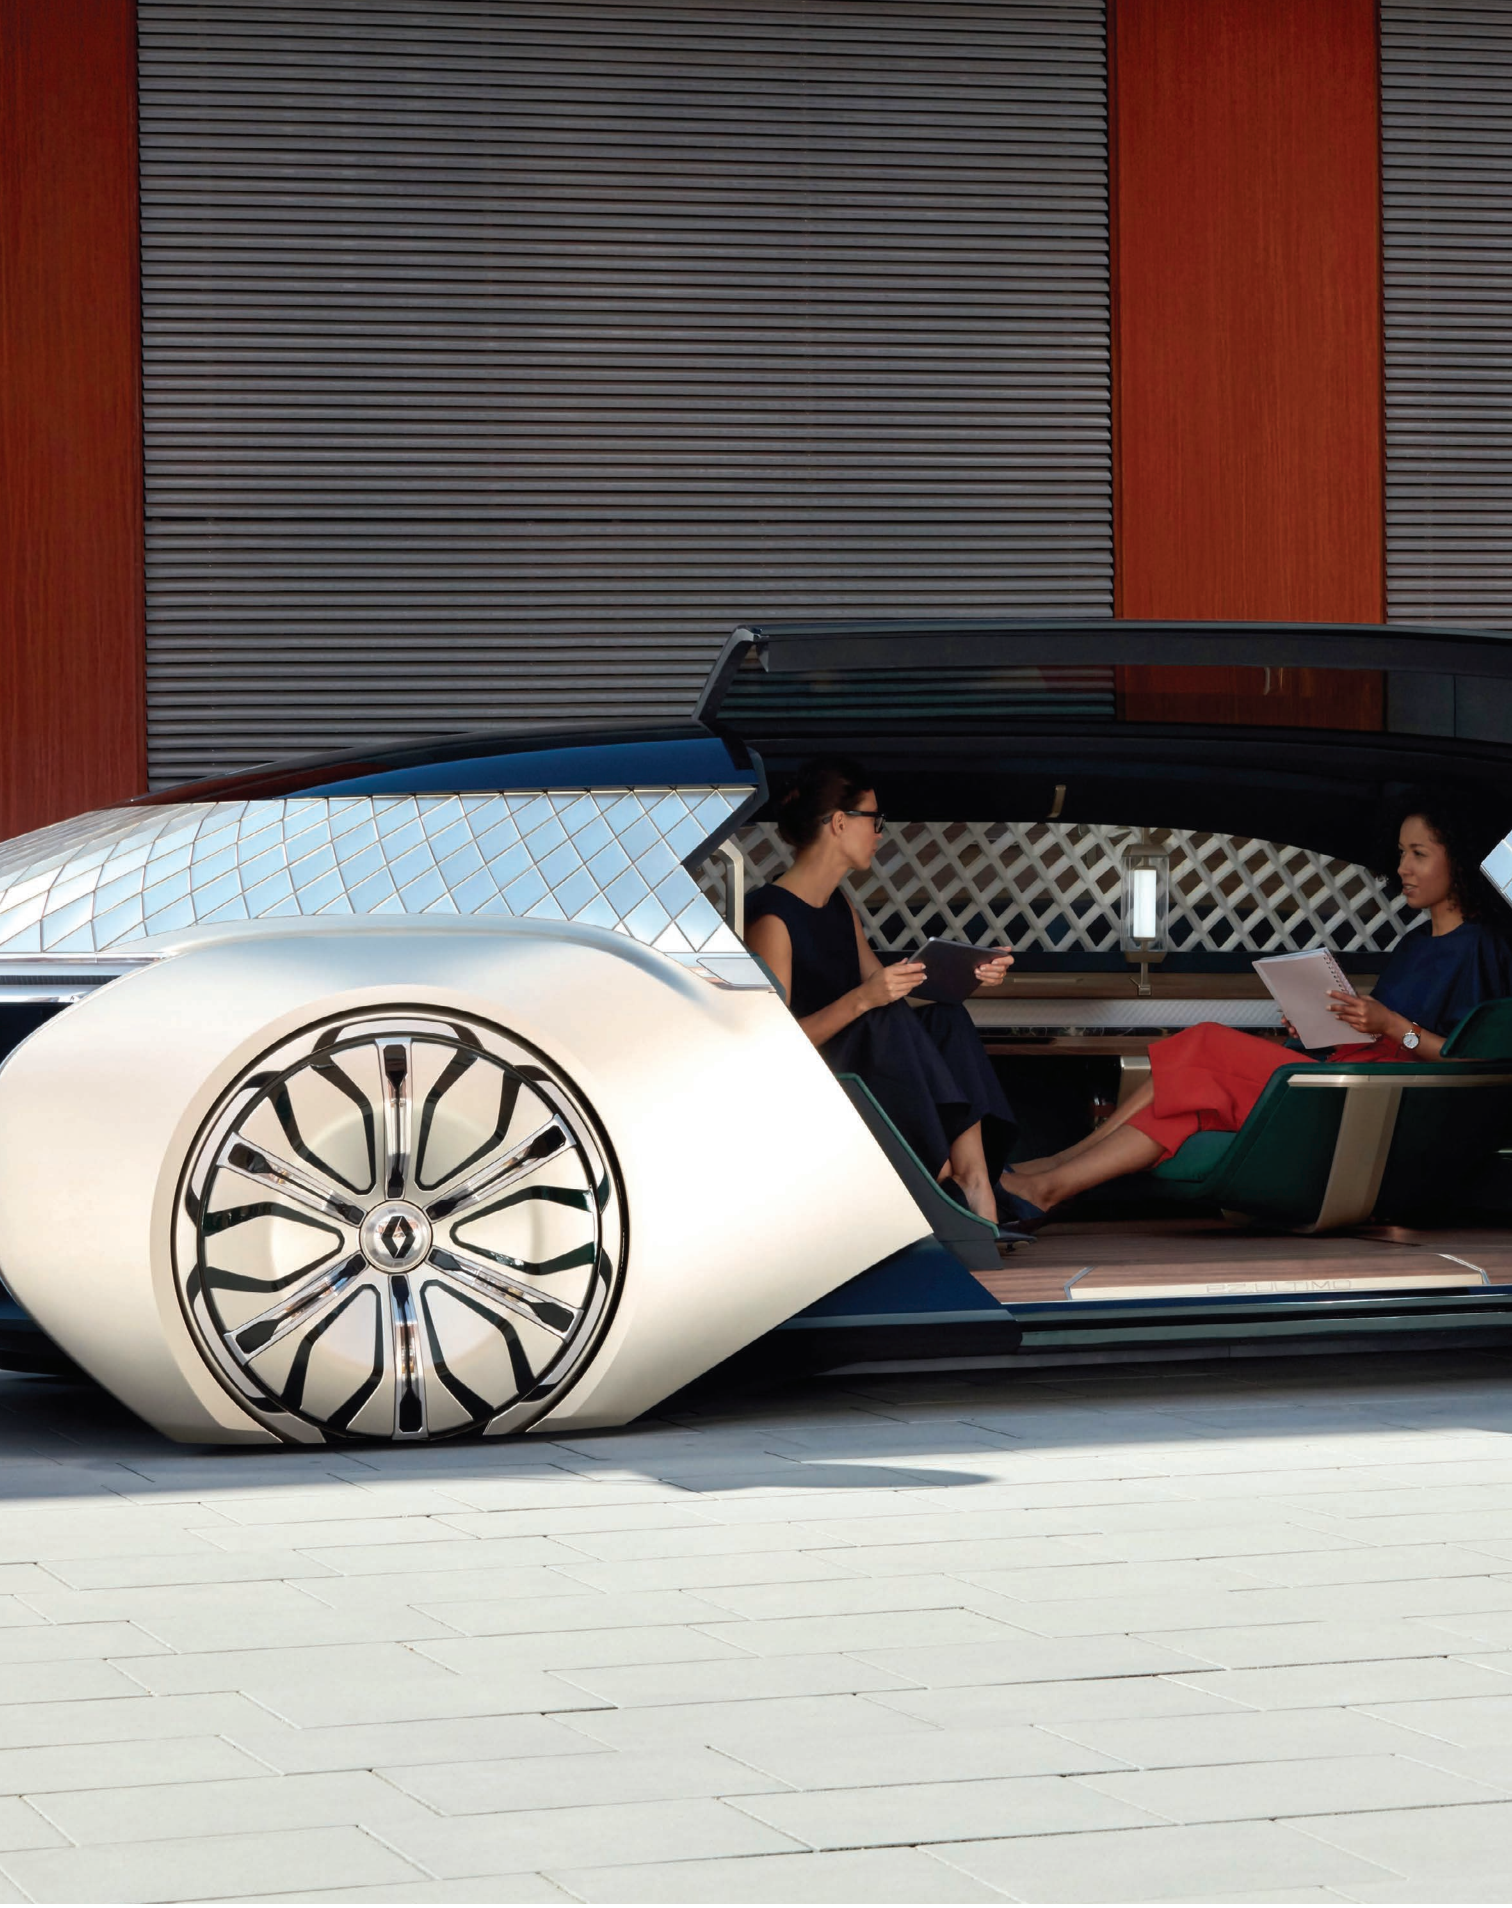

PHOTOGRAPH BY **Renault**

WHEELS FOR THE WELL-HEELED

THE ULTRARICH won't settle for an Uber to the airport. Renault's EZ-Ultimo concept car, which is essentially an executive lounge on wheels, was dreamed up with the jet set in mind. The fully autonomous all-electric vehicle is spacious, but is meant to accommodate no more than three passengers. The 600 diamond-shaped glass tiles between the body panels and the panoramic glass roof act as a giant one-way mirror that lets riders look out on the world while maintaining their privacy. EZ-Ultimo shares some traits with Renault's earlier self-driving concept vehicles—the EZ-Go for ride sharing and the EZ-Pro for delivery fleets—but the automaker is mum with regard to the particulars of the car's power train and self-driving abilities.

RESOURCES

ERICA SNYDER (3)

RESOURCES_HANDS ON

ILLUMINATING MUSICAL CODE
PROGRAM AN ELECTRONIC MUSIC PERFORMANCE IN REAL TIME

LIVE CODING IS A TYPE OF PERFORMANCE ART IN WHICH the performer creates music by programming and reprogramming a synthesizer as the composition plays. The synthesizer code is typically projected onto walls or screens for the audience to inspect as they listen to the unfolding sound. Live coding events are sometimes known as algoraves, and at these events it's common to see visualizations of the evolving music projected alongside the code. Often, these visualizations are created by a second performer manipulating graphics software in tandem with the live coder. • After attending a few algoraves in New York City (musically, the results tend to fall along a spectrum from ambient soundscapes to pounding electronic dance music, with a few detours into more experimental domains), I decided to look a little closer at the software the performers were using. I wanted to see if I could come up with my own hardware spin on creating visualizations. While I'm not yet ready to take to the stage, the results have been fun. I'd recommend that any reader interested in music or sound art should try live coding, even if they have no experience playing any traditional musical instrument. • The most popular software for live coding appears to be Sonic Pi. This is an open source project originally created by Sam Aaron for the Raspberry Pi, although it is also available for Windows and macOS. Sonic Pi's basic interface is a text editor. Apart from some performance-specific buttons, such as for starting and stopping a piece of music, it looks pretty much like any integrated development environment (IDE), in this case for a version of the Ruby language. Like Python, Ruby is an interpreted language that can run interactively. The Ruby-powered Sonic Pi IDE provides a friendly front end to the powerful SuperCollider sound-synthesis engine, which has been used for over two decades as the basis of many electronic music and acoustic research projects.

You *could* create a piece of music by typing a complete list of notes into the IDE, selecting a software-defined musical instrument plus any desired effects, such as reverb, and just having Sonic Pi play the tones. But this would eliminate the fun at the heart of live coding, which is a collaboration between the performer and the computer, in which the performer continually shapes algorithms

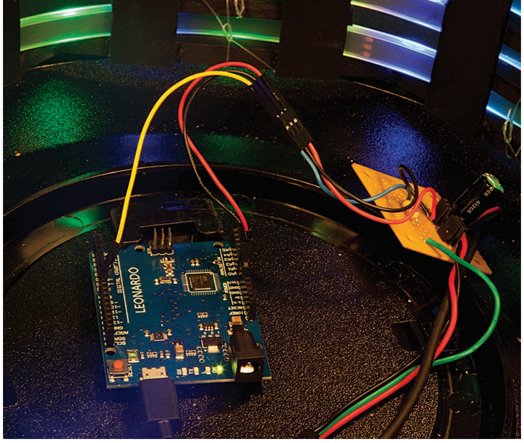

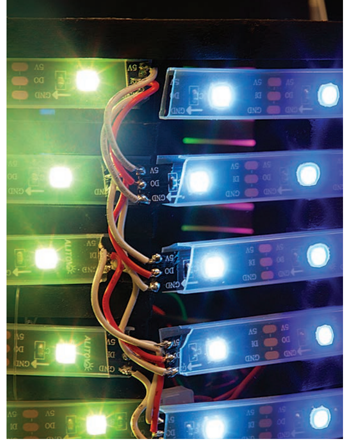

THE GREAT PRETENDER: An Arduino Leonardo [above] acts as a USB device mimicking a MIDI-enabled electronic instrument. It converts received notes into colors displayed on a strip of LEDs [right].

but leaves the work of actually determining what note to play next up to those algorithms. Sonic Pi takes care of keeping everything in sync so that the music never misses a beat.

The most recent major version of Sonic Pi introduced the ability to send and receive MIDI messages. MIDI is the venerable standard used to communicate between computers and electronic instruments. In MIDI, notes are represented by a number from 0 to 127, with notes 21 to 108 covering the range of a grand piano. Originally, MIDI required a dedicated hardware interface, but today it's quite common to see MIDI being run over USB connections.

The addition of MIDI allowed me to press-gang some hardware to visualize the music produced on the fly by Sonic Pi. A while back, I had arranged 160 programmable WS2812B RGB LEDs in five tiers, so that they act like a 32- by 5-pixel color display. I built the display on a hexagonal wooden frame and mounted it in an empty "hat box" container once used to store removable disk packs. Not only does this upcycling allow me to justify hanging onto a bulky souvenir of a bygone technology, but the roomy inside of the box allows me to hide the frame and supporting electronics, in this case an Arduino Leonardo microcontroller. The Leonardo perfectly mimics USB devices, and I've used it before to make a custom controller for a spaceflight simulator. To drive so many LEDs, I added a 10-ampere power supply, with the power and USB cables running through a small hole I cut in the box's base.

I'd already used the Arduino MIDI library, which supports MIDI over a USB, at a music hackathon where I'd converted my hat-box display into a simple light organ. I could play a MIDI file from a computer and have the display change color according to the note. But my color mapping between note values and LED colors was quick and dirty to say the least: The same color was evoked by different notes.

For my Sonic Pi visualizer, I programmed the Leonardo using the FastLED library for both performance reasons and because of its support for the HSV (hue, saturation, value) color model. Mapping a value—such as a MIDI note—to a triplet of conventional RGB values is not straightforward, especially if you want all the notes to look equally bright. In contrast, with the HSV model, it's trivial to map a note to the hue byte while keeping the saturation and value bytes fixed.

Connecting the hat-box visualizer to the Sonic Pi software was an unremarkable if fiddly voyage through the various MIDI settings on my laptop. Sending a note to be visualized does require some changes to my Sonic Pi live code, however: As each note is generated algorithmically, I capture it using an intermediate variable rather than playing it immediately in a sound-synthesis instruction as I normally would. I use the intermediate variable to send the note to the hat-box display, via the "midi_note_on" command, in addition to playing the note audibly. This allows me to program the visualizer as I program the sound code.

My next step will be to program the hat box to respond to a set of custom MIDI control commands, which will allow me to alter how notes are mapped to hue values, or even select different visualization styles, on the fly. Then you might actually find me taking to the stage. **—STEPHEN CASS**

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/sonic0919

RESOURCES_Q&A

HOW THE RASPBERRY PI INFILTRATED INDUSTRY

THE $35 COMPUTER'S CREATOR EXPLAINS WHAT DROVE THE LATEST REVISION

Seven years ago, Eben Upton created the first Raspberry Pi. As Upton told *IEEE Spectrum* in our March 2015 cover story, the Pi was inspired in part by his childhood experiments with a BBC Micro home computer: He wanted modern kids to have a simple machine that allowed for similar experimentation. Since then, the Pi has exploded in popularity, and the fourth major revision of the Pi was released in June. Upton talked with *Spectrum* senior editor Stephen Cass about the Pi 4's design, its growing commercial use, and what might be next.

Stephen Cass: How has the Pi's user base evolved?

Eben Upton: Our first year, our volume was almost entirely bought by hobbyists. But you have a lot of hobbyists who are also professional design engineers, and when their boss asked them to do something, often they used a Pi. So now you have people who are building industrial products around the Pi to resell. And then you have what we call, for want of a better word, DIY industrial, which is "I own a factory and I need control computers." And where I might have historically gone and bought an embedded PC, I'll buy a Pi. Last year we sold 6 million units and [we think as much as] half of those went to some kind of commercial use.

S.C.: How did that evolution shape the design of the Pi 4?

E.U.: We're quite lucky in that all our markets have similar requirements. The things we do to make it a better toy make it a better industrial computer. And the things you do to make it a better industrial computer make it a better hobbyist platform. For example, we had a number of people building thin-client solutions, and the feedback was that most of the people doing that wanted the option of being able to deploy two monitors. So we added that feature. Another obvious example would be the serial interfaces. Prior generations are a little underprovisioned for things like UARTs, serial ports, SPIs, I2C interfaces, but because we were redesigning the silicon from scratch, putting UARTs in was pretty straightforward.

S.C.: What other changes did you make in redesigning the Pi's system processor?

E.U.: The last three Pis have all been made using the same [40-nanometer] process, so modifications have been limited, mostly putting in a larger ARM [CPU]. If you look at the floor plan of the chip, each new ARM gets stuck on the left-hand side. The chip stretches horizontally and becomes bigger and bigger. But nothing on the right-hand side changes, and the right-hand side is where the memory controller is, where the UARTs are.... But putting in a larger ARM core means your power goes up and eventually you run out of your thermal budget.... We realized we needed to go to a 28-nm process node. And once you're going to a new process node, you might as well fix all of the wrinkles. And that's why you see two displays, many UARTs, PCI Express, Gigabit Ethernet.

S.C.: How long before the slowdown in Moore's Law affects the Pi?

E.U.: I think it's relevant to think in terms of how much is left. How much for a given thermal footprint? We've come a factor of 40 [in the Pi's computing power from the first Pi to the Pi 4]. There's not another factor of 40 left, which means you've come through more than half this process.... On some level that's intimidating. But I'm a software engineer. It's actually really nice to feel that we're going into an era where software engineering makes a contribution again! My Ph.D. is in optimizing compilers, and while I was doing it I felt really depressed, because [improving software by a factor of two was small compared to the exponential hardware improvements of Moore's Law]. But the thing about that factor of two is that it's there at the end. There will come a time when they'll be very grateful for it!

POST YOUR COMMENTS at https://spectrum.ieee.org/upton0919

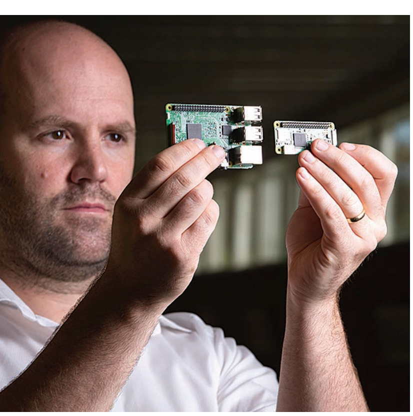

THE ARCHITECT: Eben Upton holds a Raspberry Pi and the smaller Pi Zero, designed for simpler tasks.

JEFF GILBERT/ALAMY

A-BOMB TOURISM

HANFORD'S B REACTOR IS A RELIC OF A COMPLICATED PAST

A literal pile of cylinders rises 11 meters high inside a graphite box, filling the dimensions of a cavernous hall. The towering grid of over 2,000 tubes is a jaw-dropping, neck-twisting display. Yet size and symmetry aren't all that make this a humbling sight. This is the core of a nuclear reactor, one that produced the plutonium for the Trinity atomic bomb test—and for "Fat Man," the bomb that razed Nagasaki, Japan, in 1945.

The B Reactor sits on a remote corner of the Hanford Site, a sprawling expanse in southern Washington state. Completed on 13 September 1944 after 11 months of construction, the reactor belonged to the secretive Manhattan Project, the United States–led initiative to develop nuclear weapons during World War II. Even Hanford workers were kept in the dark, learning only that their labor would support "the war effort" and yield an unnamed "product."

Today the public can visit the reactor, at least some of the time. The U.S. Department of Energy and the National Park Service offer free tours from April to November, allowing visitors to roam the halls of the world's first large-scale plutonium production complex, which shuttered operations in 1968. While the Hanford site grew to include nine reactors, only the B Reactor remains accessible—thanks largely to the former workers who fought to preserve it. Now it's part of the Manhattan Project National Historical Park, which includes sites in Los Alamos, N.M., and Oak Ridge, Tenn.

I took the tour, traveling 45 minutes by bus from the visitor's center in Richland, Wash., through the arid landscape, past puffs of sagebrush and bright green potato plants. Rattlesnake Mountain looms above and the Columbia River courses by, replenishing the crops as it once cooled the nuclear reactors. A passenger asks how much radioactive contamination we'll be exposed to at the site, to which the guide replies, "None."

The museum tour starts with the showstopper: the reactor core. The 2,004 aluminum tubes stretch back about 13 meters, though we can see only their front ends. During production, physicists took 22-centimeter-long slugs of naturally occurring uranium and slid them into the horizontal tubes. Surrounded by graphite, the uranium was transformed by nuclear fission into plutonium. Treated river water pumped through the reactor at a rate of 280,000 liters per minute.

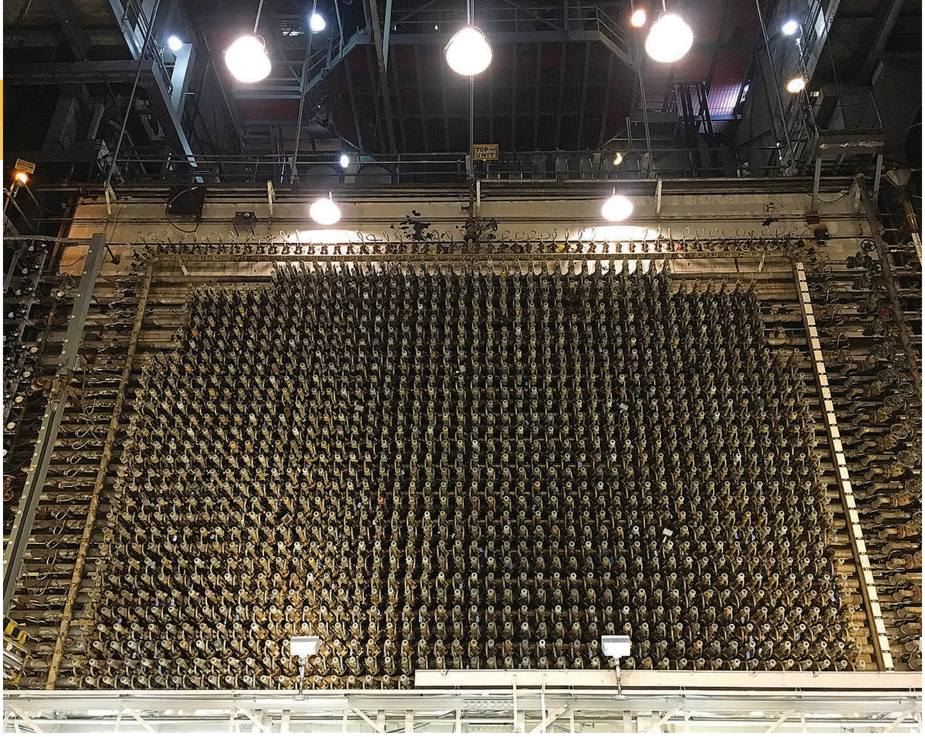

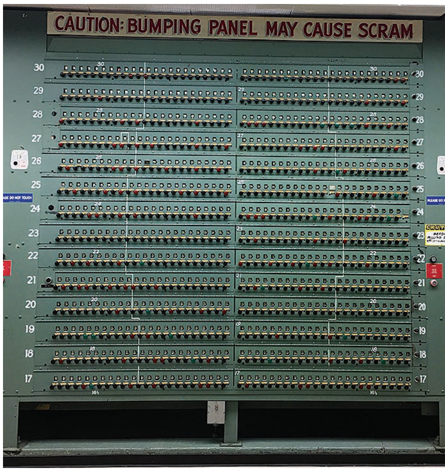

THE PLUTONIUM FACTORY: [Clockwise from top left] The front face of the B reactor is made from thousands of aluminum tubes. Although the site is decommissioned, much of the original signage is still in place, such as these signs in front of a distribution panel and above a control board that monitors water pressure inside the reactor's tubes. While safe to visit, the site is still monitored for radiation through devices such as wall-mounted dosimeters. The valve room emphasizes the industrial nature of this reactor, where plutonium for the United States' atomic bomb arsenal was made.

MARIA GALLUCCI (5)

To the side, in a narrow control room, thousands of tiny knobs cover the wall like red, black, and green polka dots. The mechanical gauges monitored water pressure inside each of the tubes. Engineers often worried that bumping a gauge would throw off the readings and cause an emergency "scram," or the rapid insertion of 29 vertical safety rods to shut down the reactor. (It never happened.) Opposite the reactor, a gaping pit of metal pipes and manual valves shows how cooling water snaked into the pile.

Hundreds of original artifacts are sprinkled throughout the building. Rubber masks and thick gloves sit in a former decontamination room. Rotary phones and typewriters adorn wooden desks. Walls are covered in long-outdated calendars and hand-painted warning signs. It is a shrine to World War II–era ingenuity.

Yet the B Reactor tour doesn't dwell on the consequences of the Manhattan project. Visitors hear little of how, when the United States bombed Hiroshima and Nagasaki in August 1945, hundreds of thousands of people died from direct impact and the lingering effects of radioactive fallout. And the contaminated soil, groundwater, and leaking tanks of waste just south of the reactor are framed as necessary by-products of national security.

Robert Franklin, president of the B Reactor Museum Association, acknowledges that the site is "a real Pandora's box of history." He says the tour does encourage visitors to engage in complex and difficult discussions about U.S. nuclear history.

"On the one hand, it's a marvel of science and engineering," he says of the B Reactor. "But on the other hand, it helped to create this world of uncertainty, fear, and anxiety."

—MARIA GALLUCCI

POST YOUR COMMENTS at https://spectrum.ieee.org/hanford0919

3 WAYS TO LEVERAGE YOUR SOCIAL MEDIA

BOOST YOUR CAREER WITH TWITTER

Scientists and engineers are increasingly using social media platforms such as Twitter and Facebook to promote their work. Look at Bill Nye. The mechanical engineer, TV and podcast host, and CEO of the Planetary Society, Nye is a Twitter superstar with 5.9 million followers. But can social media benefit regular engineers and academics as well? Yes, it can–once you hit a 1,000-follower threshold on the microblogging site, according to a June 2018 study in the Canadian journal *Facets*. That's when your tweets start to reach a broader audience outside your immediate fraternity. If you are thinking of navigating the world of social media sites, here are some ways it can be useful and suggestions to get you started.

Keeping the public informed: While there are plenty of social media platforms, such as Facebook, Snapchat, and Instagram, to choose from, Twitter seems to be the favorite of the research community and so attracts members of the general public interested in such topics.

Consequently, Twitter can help you counteract fake news in your discipline. Thomas G. Dietterich, an emeritus professor at Oregon State University, specializes in artificial intelligence and machine learning. He has 16,500 followers on Twitter and uses the social media channel for fact-checking misinformation floating around the Internet about AI. "Both media companies and research institutions have incentives to exaggerate the generality and significance of research," says Dietterich. "If the public does not understand the technology, it may oppose it in situations where it could have great benefit or enthusiastically support it in cases where it could be dangerous."

You can also use social networking sites like Twitter and Facebook to publicize your research and achievements. If that's your goal, you should focus on boosting the number of your followers. "You can do that by joining conversations led by people who have many followers," says Dietterich. "You also gain followers by posting interesting pieces (e.g., on a blog) and then linking to them from Twitter." | CONTINUED ON PAGE 50

DIGITAL SNAKE OIL

AS MORE TECHNOLOGY FIRMS PRODUCE wearables, apps and connected medical devices that claim to help people live better or treat diseases, we need to draw a line between digital wellness and digital medicine. The entire health care industry needs to implement rigorous standards that can help differentiate between truly therapeutic products and the digital equivalent of snake oil. • Today, consumers and doctors are bombarded with claims. Apple says the Apple Watch can detect if the person wearing it is going into atrial fibrillation. Researchers believe they have developed an app that can tell if you are depressed simply by monitoring how you type and interact with the screen. Companies are pushing home versions of medical devices for detecting respiratory diseases in children, spotting chronic obstructive pulmonary disease, and testing urine that all claim to deliver clinical data. • In the United States, the Food and Drug Administration has clamped down on a few offending apps, such as several that purport to diagnose concussions. But there are plenty of devices that straddle the line between digital wellness and actual digital medicine. • Instead, those defining the distinction are pharmaceutical companies and the startups working with them to improve clinical trials. They're creating benchmarks, writing best practices for device security, and urging doctors to adopt only technology that can clearly meet these standards. Maybe most important, they're using devices that meet these standards in clinical trials to make the trials more inclusive and more efficient. With luck, their efforts will make digital medicine less hype-driven and more data-driven instead. • Andy Coravos, CEO of Elektra Labs, in Boston, has built a startup that tests qualities of today's connected devices and creates what you can think of as an easy-to-understand label for each device that lets pharma companies know how well different products fare. The company's benchmarks verify device claims such as whether a Fitbit's accelerometer actually can measure 100 steps. The label also validates claims against a clinical goal—in other words, does that 100-step measurement actually mean anything?

The label indicates a device's usability and security, and the device maker's data-retention practices. Security and data retention are especially important: Coravos says that for the first time, pharma companies have to consider whether or not a medical device could be hacked, and where liability for that might lie.

Monarc Bionetworks, in Redwood City, Calif., is another startup trying to codify processes so that consumer devices can be used in clinical trials. Monarc is working with wearable-device companies, providers of electronic health records, and pharmaceutical companies to make sure the data that gets used is vetted according to clinical standards.

Bringing all of these efforts together is the Clinical Trials Transformation Initiative (CTTI), a public-private partnership cofounded by Duke University and the FDA. In February, it issued 17 recommendations to guide companies and clinicians in using smartphones and mobile tech to improve clinical-trial quality and efficiency.

The recommendations include establishing benchmarks like those used by Elektra Labs as well as how to introduce them in new tech practically. For example, CTTI recommends that a study ensure access to tech help for patients and clinicians, and suggests regular conversations about data use and storage.

This is a brave new world in which medicine can truly become personalized. But to give digital medicine a healthy future, we have to ensure that what's on the market is actually therapeutic and not just electronics wrapped in promises. ■

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/digitalsnakeoil0919

ILLUSTRATION BY Jude Buffum

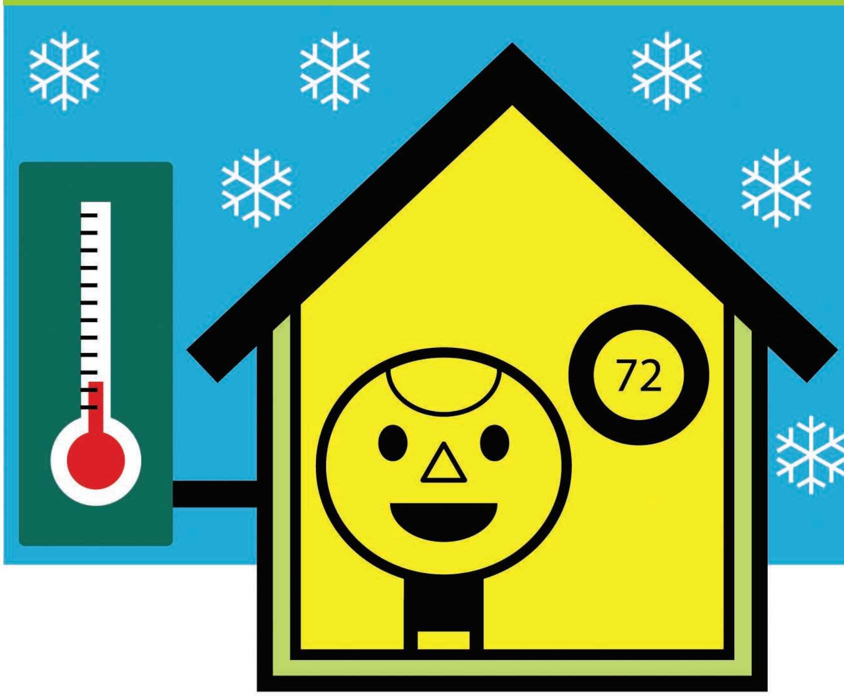

BRICKS AND BATTS

FIRST IMPRESSIONS OFTEN LEAD TO wrong conclusions. I well remember receiving a friendly welcome at the residence of a European ambassador in Ottawa and, in the very next sentence, being told that this house was perfect to withstand Canadian winters because it was made of real brick and stone–not like those flimsy North American wooden things, with hollow walls. My hosts then swiftly moved to other matters and, in any case, I did not have the heart to belittle the insulating qualities of their handsome home. ● The error is easy to understand, but mass and density are better indicators of sturdiness than of insulating capability. A brick wall obviously looks more substantial and protective than a wall framed with narrow wooden studs and covered on the outside with a sheet of thin plywood and aluminum siding and on the inside with vulnerable gypsum drywall. Angry European men do not make holes in brick walls. ● Decades ago, when oil sold for US $2 a barrel, most pre-1960 North American houses usually had nothing more to keep out the cold than the air space between the plywood and drywall. Sometimes the space was filled with wood shavings or shredded paper. Yet, remarkably, even that feeble combination provided a bit more insulation than solid brick. ● The insulating value, or thermal resistance, is measured in terms of R-value. It depends not only on the composition, thickness, and density of the insulation but also on the outdoor temperature and moisture. A framed wall from 1960 had roughly the following R values: aluminum siding (0.6), thin plywood (0.5), air space (0.9) and drywall (0.5). It all adds up to only about 2.5. Standard brick (0.8) plastered on both sides offered no more than 1.0. Hence even a flimsy 1960 mass-built North American wall insulated at least twice as well as did Europe's plastered brick. ● Once energy prices began to rise and more rational building codes came into effect in North America, it became compulsory to incorporate plastic barriers and fiberglass batts–pillowlike rolls that can be packed between the wooden frames, or studs. Higher overall R-values were easily achieved by using wider studs (two-by-six) or, better yet, by double-studding, which involves building a sandwich from two frames, each one filled with insulation. (In North America, a softwood "two-by-six" is actually 1.5 by 5.5 inches, or 38 by 140 millimeters.) For a well-built North American wall this means adding insulation values of drywall (0.5), polyethylene vapor barrier (0.8), fiberglass batts (20), fiberboard sheathing (1.3), plastic house wrap (Tyvek ThermaWrap at 5) and beveled wood cladding (0.8). Adding the insulating value of interior air film brings the total R-value to about 29.

Brick walls, too, got better. To keep a desired outer look of colored brick, an old wall can be retrofitted from the inside by putting wooden battens (thin strips that hold insulation in place) on the interior plaster and attaching insulation-backed gypsum board integrated with a vapor membrane to keep out moisture. With 2-inch insulated plasterboard, this will triple the previous overall R-value, but even so, the insulated old brick wall will remain an order of magnitude behind the two-by-six framed North American wall. Even people who are generally aware of R-values do not expect to see such a large difference.

However, all this wall insulation can reach its potential only if the windows don't bleed heat. Today's best tripled-paned windows, filled with argon and topped with low-emissivity coatings, provide an R-value of up to 7.5. That's worse than a good wall but vastly better than the R 0.9 of an old single pane held in (draftily) by disintegrating putty. Keeping warm is both an art and a science. ■

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/insulation0919

ILLUSTRATION BY J.D. King

THE EXPIRATION DATE OF KNOWLEDGE

I WAS LISTENING RECENTLY TO SOME engineering graduates talking about their current research efforts. As I was forcibly immersed in the minutiae of opaque mathematics, the thought came to me that this was really difficult work for potentially small gains. But on the heels of that came another thought: These engineers were really skilled. I was greatly impressed by their depth of knowledge and facility with analysis. How did they learn so much in their few years of university training? After all, there is so much more to know now than there used to be, and every year it gets ever more overwhelming. • One explanation is that, as new knowledge accumulates, some old knowledge becomes irrelevant and falls off the knowledge stack. Almost all the college course work I took long ago is now useless in itself, although what remains is an engineering mind-set and a mathematical grounding. Perhaps every course should have a sell-by date. Indeed, in retrospect I now realize that a number of the courses I took were already well past their sell-by dates when I took them. I remember too when the technical library in the lab where I worked was shut down and all the books were offered free to the staff. Almost all of them went unclaimed; no one wanted them. The problem is that we never quite know *when* a particular course or book will become obsolete. • But the purging of obsolete knowledge is probably insufficient in itself to make room for the new stuff, as there seems to be an exponential increase in knowledge. The complexity of our work is always increasing, similar to the increase in entropy decreed by the second law of thermodynamics. For many decades we've been driven by Moore's Law, which has urged us to embrace and exploit complexities that were unthinkable in previous decades. Even as Moore's Law wanes, I feel sure that the general law of exponential increase will continue the trend.

New engineers often enter fields that have become well plowed, and so find themselves pushing against physical and theoretical limitations. The issues are complex, and the incremental gains may be small. I imagine an index of measurement–potential gain divided by complexity of required work. Normally, this index grows ever smaller, but fortunately new fields and new techniques open up periodically, and engineers rush in to take advantage. For example, I frequently hear about engineers employing machine learning in some new and creative way.

As the new engineers come out of school, they are also empowered by the continual rise of new tools. They work with networked computers, whose software embeds the knowledge and techniques of analysis and design accumulated by others. I remember when we used to wander down the halls asking other engineers how to solve some problem. Now we put our questions to the computer, and by doing so, ask all the other engineers and scientists of the networked world. You don't have to know everything yourself, only how to ask the questions.

So the challenge is great, but these graduates are up to it. I'm thinking, however, that there must be some value in the experience of older engineers. But that is a subject for another time; this is a salute to the new guard. ■

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/reflections0919

ILLUSTRATION BY Andrea Mongia

THE MAGLEV HEART

The decades-old quest for an artificial heart takes a surprising turn

By Nicholas Greatrex, Matthias Kleinheyer, Frank Nestler & Daniel Timms

Photography by Peter Adams

THE INVENTOR: Daniel Timms, founder of Bivacor, holds the 650-gram machine that could save the lives of men, women, and children with heart failure.

FOR MORE THAN 50 YEARS, cardiac surgeons and biomedical engineers at the Texas Heart Institute (THI) have been questing for an artificial heart that can fully replace natural ones, which are in terribly short supply for transplant. They've seen their share of metal and plastic contraptions that used a variety of pumping mechanisms, but none of these machines could match the astounding performance of the human heart.

In April 2019, the possible culmination of that long quest was inside a shaggy brown cow, which stood peacefully chewing its cud at a THI research facility in Houston. The animal was part of a 90-day trial in which it lived its life powered by an implanted artificial heart made by our company, Bivacor. Throughout the trial, the calf stayed healthy and energetic, and gained weight at a normal rate. It even jogged on a treadmill for 30-minute stretches.

Our company is now working toward human trials of our device. It relies on a dramatic new approach: Rather than using a mechanical pump that mimics the structure and actions of the four-chambered human heart, it uses a spinning disk, suspended in a magnetic field. With just one moving part, the Bivacor heart is able to send oxygen-rich blood out to the body and return oxygen-depleted blood to the lungs.

We had to overcome many technical challenges to make an artificial heart that's small, biocompatible, energy efficient, and durable. Consider that the human heart beats about 112,000 times a day, which adds up to 42 million times a year, and you'll understand the magnitude of the challenge.

We've tested the Bivacor heart in 15 cows so far. While the need for animal testing is unfortunate, it's the only way to prove the device's safety and move forward to clinical trials in humans. These Corriente calves, which are rela-

PHOTOGRAPH BY Peter Adams

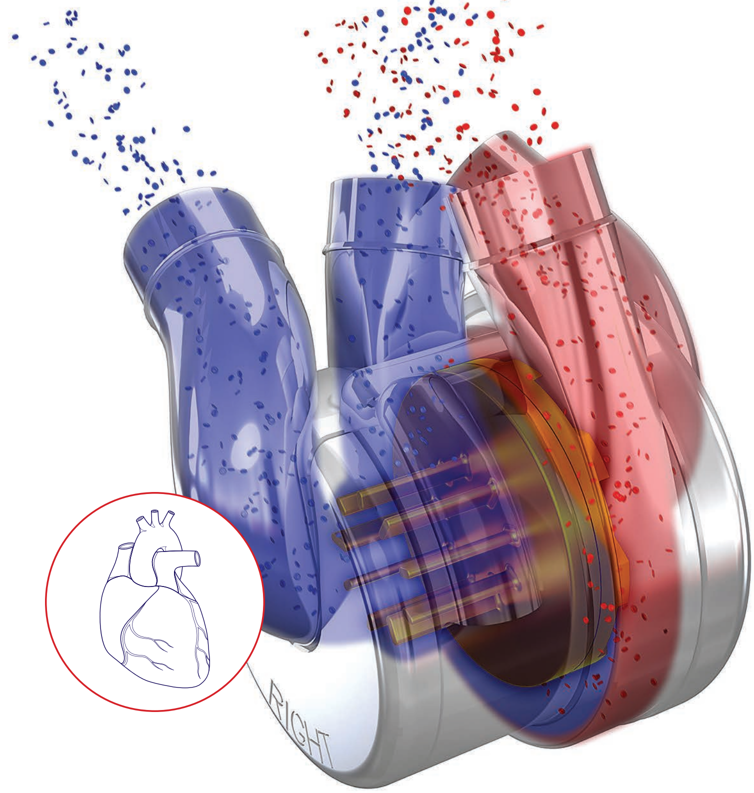

tively small, are the right size to serve as analogues for adult patients. We've also implanted the Bivacor heart in several sheep, which are more representative of patients with smaller bodies, including children. Our tests have shown that the heart holds up well: With its one moving part levitating in a magnetic field, there's no worry that friction and mechanical wear will cause the machine to give out. Our tests have also shown that the device can adapt to the user's cardiovascular requirements.

One of us, Bivacor founder Timms, began working on this project 18 years ago. [For more about his personal story, see "Taking Lessons to Heart," p. 2.] Through the years, the central concept of the maglev heart has stayed the same, though our engineering team has greatly refined the technology. We're thrilled that the device is nearly ready for human trials because we think the Bivacor heart is the solution that severely ill heart-failure patients have long been waiting for.

IN HEART FAILURE, the heart becomes unable to pump enough blood to keep the body healthy and strong. At least 26 million people around the world are living with the disease, and the number is rising as populations age. Patients with severe heart failure have a bleak outlook: Their best option is a heart transplant, but the limited number of donor hearts means that only about 5,000 patients around the world receive transplants each year. Thousands more patients are eligible for transplants, and some die while waiting for a donor organ.

Cardiologists have long dreamed of a mechanical replacement. In 1969, THI physician Denton Cooley implanted the first "total artificial heart" in a patient who was awaiting a transplant, keeping him alive for the 64 hours it took for a donor organ to arrive. However, that patient died shortly after the transplant surgery from an infection, and Cooley's

HEART TECH: In Bivacor's engineering lab, founder and CEO/CTO Daniel Timms tinkers with the parts of the latest prototype of the artificial heart [above left]. Its centrifugal pump [above right] uses a single rotating disk, which is suspended in its casing via magnetic levitation. The left side of the pump stands in for the anatomical heart's left ventricle, doing the job of sending oxygenated blood [red] out to the farthest reaches of the body; it has wide impeller vanes that produce enough pressure for the task. The right side of the pump substitutes for the heart's right ventricle; its narrower impeller vanes return deoxygenated blood [blue] to the lungs. The Bivacor device doesn't need components to stand in for the heart's left and right atrium, which take in blood and send it to the ventricles.

LEFT: ISTOCKPHOTO; RIGHT: BIVACOR

team shelved the device due to concerns about reliability and compatibility with the human body.

Since then, a handful of other total artificial hearts have been developed, and a few have made it into human trials. But these devices were large, heavy, and prone to mechanical failure, and only two gained regulatory approval in the United States. The SynCardia Systems artificial heart was approved in 2004 as a "bridge to transplant" and is currently being tested as a permanent replacement. However, users must carry a 6-kilogram case containing a loud air compressor that's attached to tubes penetrating the abdomen, where the air drives the device's two pneumatic pumps. A second artificial heart, the AbioCor, gained approval in 2006, but it was discontinued almost immediately, when the company behind the device decided it wasn't commercially viable.

The most common mechanical support now in use is what's called a left ventricular assist device (LVAD). This machine adds its pumping power to an ailing heart, focusing on the left ventricle, which pumps oxygen-rich blood through branching arteries that reach up to the brain and down to the toes. LVADs are now being used as both temporary aids for patients awaiting transplants and, in cases where transplantation isn't an option, as permanent additions to patients' chest cavities. However, LVADs can cause the right ventricle to falter, requiring intensive drug treatments and sometimes the implantation of a right ventricular assist device.

At Bivacor, we wanted a true long-term mechanical replacement for the entire heart. Unlike most earlier attempts, our effort didn't set out to mimic the heart's natural pulsatile pumping mechanism, with valves that open and close as the left ventricle pushes blood out to the body and the right ventricle pushes blood to the lungs. Most prior artificial hearts used positive displacement pumps to achieve this effect: Their two artificial chambers were bifurcated with membranes that flexed forward to push blood out through mechanical valves.

We opted instead for a centrifugal pump that propels a continuous stream of blood into the arteries. Such a pump has no valves, so a patient using the Bivacor heart in its most basic mode would have no pulse. But we recently adapted our device to give it a pulsatile outflow option, as we'll explain below. We want clinicians to look at the cardiac monitors for patients with implanted Bivacor hearts and see the familiar rhythmic readouts they've seen since medical school.

THE BIVACOR HEART would fit in the palm of your hand—it's about 650 grams, slightly heavier than an adult human heart. Its shell is made of titanium, a noncorroding material that almost never triggers an immune response. Patients will wear a 4-kg external controller pack that contains two rechargeable batteries (providing about 5 hours of

IN THE LOOP: To rapidly test their artificial heart prototypes, and to reduce the need for animal testing, the Bivacor team built an apparatus that replicates the human circulatory system [above], simulating the flow of blood through the body and the lungs. By connecting a prototype to the testing rig, Bivacor engineers can watch how it propels a synthetic bloodlike liquid through the tubes and valves. The engineers test a prototype's sturdiness [opposite, left and middle] by keeping it hooked up to a special durability rig for several months. The device's casing [opposite, right] is made of titanium, which is biocompatible. The parts that look golden have a coating of titanium nitride.

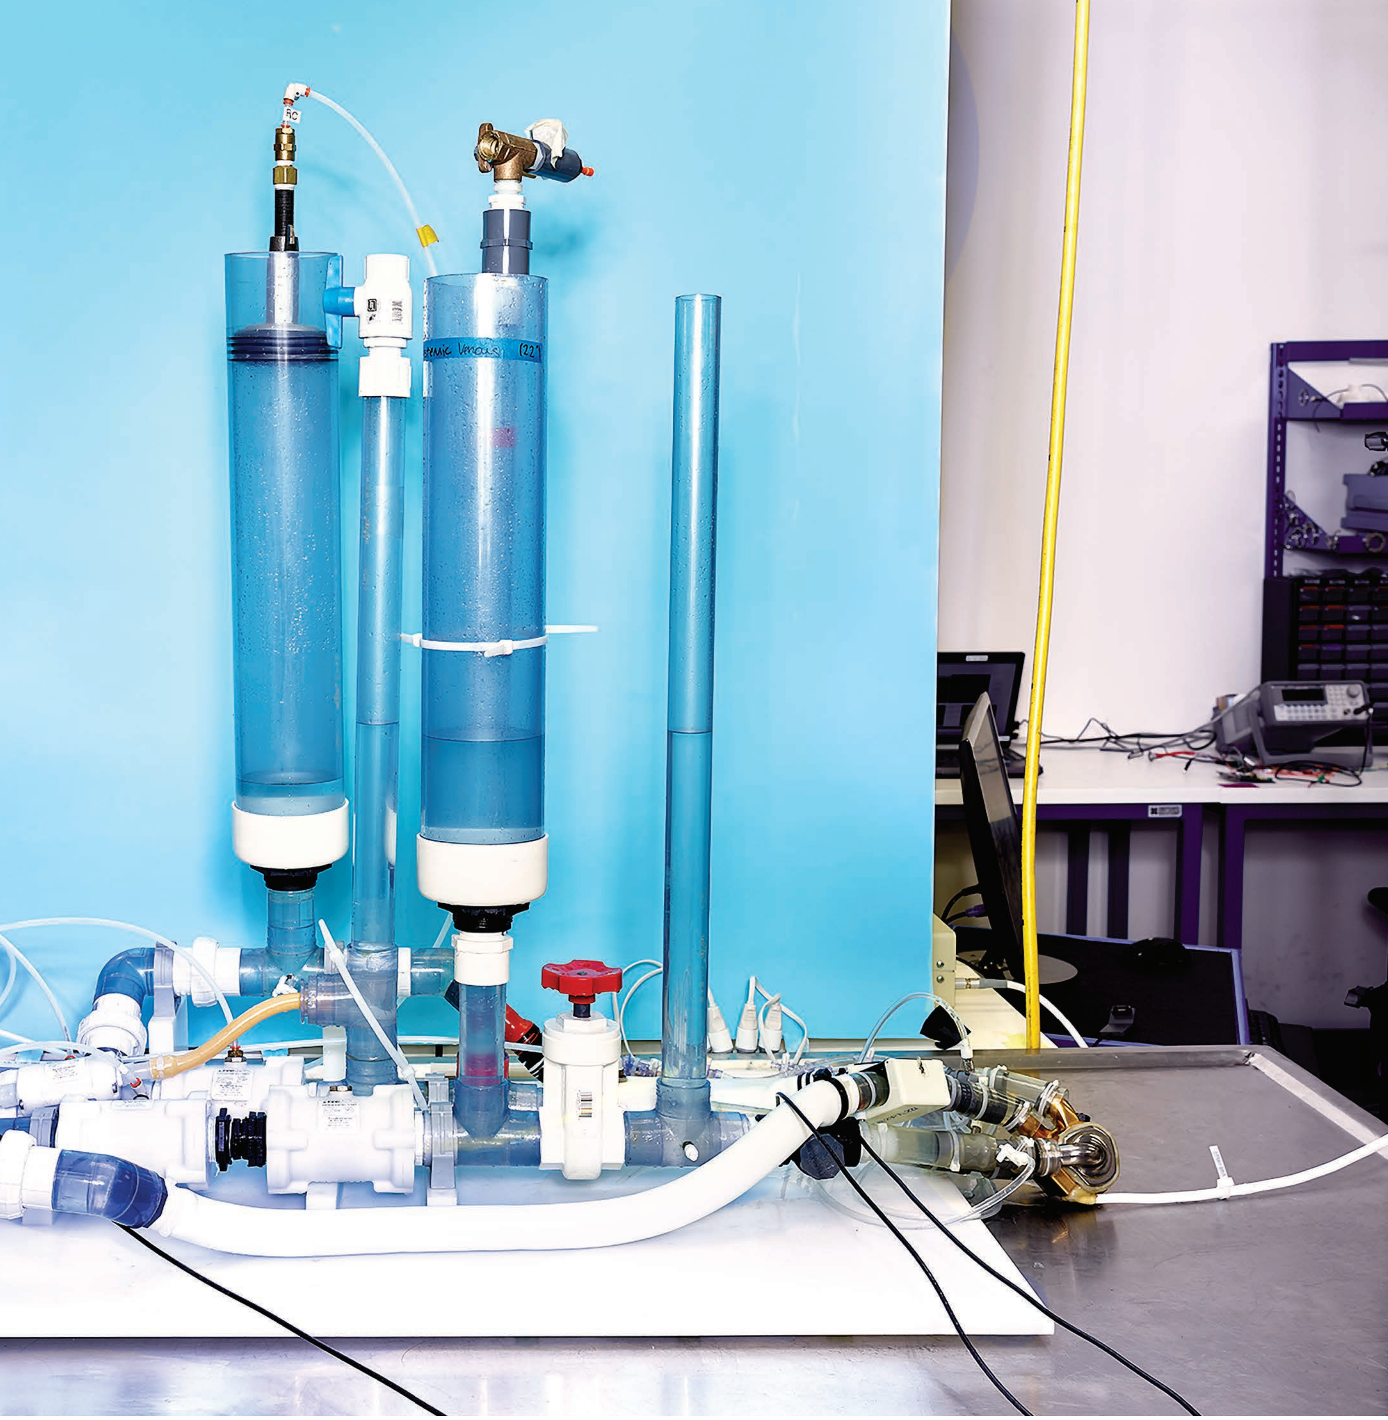

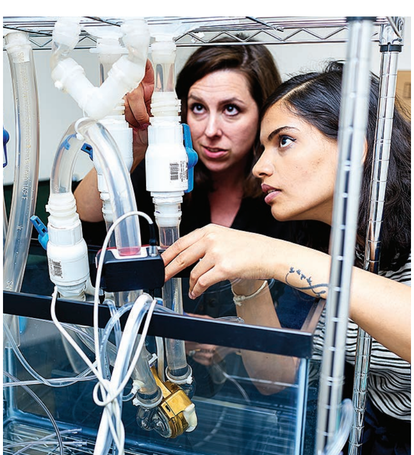

PHOTOGRAPHS BY Peter Adams

EVOLUTION OF AN IDEA: Timms came up with the original concept of an artificial heart with a centrifugal pump in 2001 and has been developing the technology over the past 18 years. The first concept model [left] tested the hydraulic feasibility of the design, with a single rotor pumping fluid in two directions. Later prototypes [middle left and right] were used for the crucial tests in cows. The final design [right] is more tailored for human anatomy.

operation each), although they can also plug in directly to a power outlet.

Throughout our design process, we used 3D printers to make both titanium and plastic parts for our prototypes, allowing us to rapidly experiment with different geometries. For testing, we built a hardware simulation of the human circulatory system in our engineering office in Los Angeles; this mock-up allows us to validate a device's function thoroughly and repeatedly in a controlled environment, and reduces the need for animal testing.

The primary design innovation in the Bivacor heart is its simple construction, with one motor and one rotating disk that simultaneously supports the pumping of blood to both the body and the lungs. The rotating disk is completely suspended in a magnetic field, operating under the same "maglev" principle that has been used by high-speed trains. The disk has open impeller vanes on both sides, one larger set that pumps blood at the high pressure necessary to send blood throughout the body, and another smaller set that pumps blood at lower pressure to the lungs. Each side of the disk can pump more than 12 liters per minute, more than enough output for patients engaged in moderate exercise.

While some blood from the two sides mixes around the edges of the suspended disk, this isn't a concern, because of the direction of the flow. Some oxygenated blood leaks from the high-pressure side to the low-pressure side, which means that some already-oxygenated blood returns to the lungs. And this leakage is actually a design feature, not a bug. The wash of blood around the disk cleans out the casing and ensures that there are no areas where stagnant blood can form into dangerous clots.

The motor's stator provides rotational torque by coupling to a set of permanent magnets in the rotor disk. During normal operation, it spins the rotor at speeds of between 1,600 and 2,700 rotations per minute. The attractive forces between the motor and the rotor are counteracted by the magnetic bearing on the opposite side of the rotor, which actively controls the rotor's position within the casing. This active control system is necessary because the rotor naturally moves around as the patient walks, climbs stairs, jumps, and generally goes about daily activities. It's important to keep the rotor properly suspended and to prevent it from bumping into the sides of the casing, which could damage the components and smash blood cells.

This positional control system works as follows: Tiny contactless sensors send out magnetic fields that interact with the rotor, determining its exact location many times per second. If the rotor is moving in one direction or another, the control system puts electrical energy into electromagnetic coils within several actuators, causing them to cancel out that movement.

The design of the magnetic bearing actuators was critical, particularly because they had to be small, light, and energy efficient, yet strong enough to compensate for all the jiggles and joggles created by a person on the move. We used computer simulations of the magnetic field to optimize the design, experimenting with different materials and geometries to find the configuration that would provide sufficient force within a small space. Keeping our device small and light means that it will fit inside people with small bodies, including women and children.

To improve the efficiency of our artificial heart, we also integrated a "zero power" controller into the suspension system. This controller monitors the additional electromagnetic power used by the bearing as it responds to external forces acting on the rotor, then moves the rotor to a position where the permanent magnets of the magnetic bearing system can provide a balancing force. This system doesn't produce instant adjustments—that's the job done by the main stability controller—but it does reduce the amount of power used by the magnetic bearing when exposed to external forces.

A unique feature in the Bivacor heart is that the rotor can shift along its axis of rotation to change the amount of blood moved by the left and right sides of the pump. When the rotor moves toward the left side of the casing, it brings the impeller vanes closer to the casing wall. In this narrow space, most of the blood is whirled around by the vanes and little flows over the vanes' tips, thus increasing the left pump's efficiency and consequently its output of blood to the body. This feature is useful for quick adjustments necessary in transitions, such as when a patient stands up.

PHOTOGRAPHS BY Peter Adams

Our device is elegant in its cooperation with the human body, adapting automatically to a person's activity levels. When the patient is exercising, the mechanical heart will pump more blood out to the muscles, just as a biological heart does. This adjustment is accomplished by attending to the body's feedback: When a person starts to run on a treadmill, the hardworking muscles in the legs and elsewhere use up oxygenated blood faster and squeeze the deoxygenated blood back into the circulatory system. That increased blood flow into the Bivacor heart causes both sides of the pump to move more blood, without requiring an increase in rotor speed. This seemingly basic functionality is actually the result of careful optimization of the hydraulic design. We relied on our 3D printer to make numerous experimental prototypes with casings and impeller vanes of slightly different shapes.

Biocompatibility is one of the biggest challenges in our field, because the interactions between a mechanical device and biological systems are so complex. For example, the delicate blood cells and other blood components can be damaged by rough transit through a device. As a key design feature of the Bivacor heart, we ensured that the blood has plenty of clearance between the levitating rotor and the casing and conduits. All flow paths have clearance gaps of at least 240 micrometers during normal operation, which is more than 20 times the size of a red blood cell. This design reduces the shear forces to which the blood is exposed, and also ensures that there's no blood stagnation within the casing.

The pump can run at constant speed, producing continuous blood flow at a constant pressure, and in our early experiments we concentrated on testing this "pulseless" mode. But it's easy to change its speed, and our later experiments proved that controlled speed changes could produce a wide range of flow and pressure characteristics. Running the pump first at high speed (sending out more blood) and then at low speed (sending out less) creates something resembling a biological heart's pulse; rapidly alternating these two speeds creates something that looks like a normal heartbeat.

We're now working primarily in that pulsatile mode. Within cardiology, there's an open debate about whether a pulse is necessary for good health. Some patients with implanted LVADs that produce continuous (and thus pulseless) flow experience medical problems such as gastrointestinal bleeding, yet it's not clear whether their devices are the cause. We hope that our device, which can be used in either continuous flow or pulsatile mode, can contribute to the scientific investigation of this important topic.

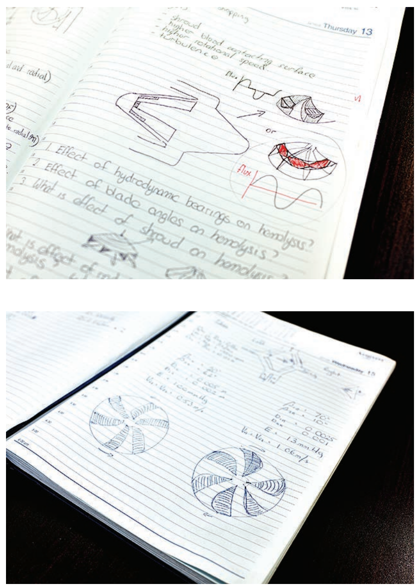

SCRATCH-PAD BEGINNINGS: Timms first sketched out the basic idea that became the Bivacor artificial heart in a notebook while he was getting his Ph.D. in biomedical engineering—shortly after his father was diagnosed with heart failure.

THE BIVACOR HEART has undergone numerous iterations since coauthor Timms began working on the original concept 18 years ago. Timms began the project in his native Australia, and since those days, the development team has benefitted from a network of passionate collaborators in Germany, Japan, Taiwan, and the United States. During various stages of the project, the team relocated to international laboratories to make use of our collaborators' expertise. Currently, our team is based in the United States (in Los Angeles and Houston) and Australia (in Melbourne and Brisbane), and we're focused on turning our prototype into a commercial product.

We're quite satisfied with the design of our device. Now we're standardizing the production process: Our 3D-printing techniques were well suited for the development phase, when we wanted to test prototypes and iterate rapidly, but now we're switching over to precision machining, which will give the device's parts a smoother polish and more exact and consistent dimensions. These procedural changes will make larger-scale manufacturing possible. We're also documenting every step of both the production and implantation process to prepare for our first clinical trials.

By the end of 2019, we'll have a clinical-grade system. Then we'll submit our request to the U.S. Food and Drug Administration for an early feasibility study in human patients, which we hope will commence in 2020.

We are already envisioning these first human trials. Gravely ill patients will go into the operating room with failing biological hearts beating feebly in their chests, and come out of surgery with smoothly functioning Bivacor artificial hearts whirring away. As the devices send powerful streams of oxygen-rich blood coursing through their bodies, we hope the patients will quickly find their strength returning. If these patients can rise from their hospital beds, hug their family members, and continue their lives for many years to come, we'll have taken a great step forward in the long quest for a total artificial heart. ■

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/bivacorheart0919

PHOTOGRAPHS BY Peter Adams


SUPERCHARGED

PORSCHE'S FAST-CHARGE POWER PLAY

The new, all-electric Taycan will come with a mighty thirst. This charging technology will slake it

BY JOHN VOELCKER

ILLUSTRATION BY MCKIBILLO

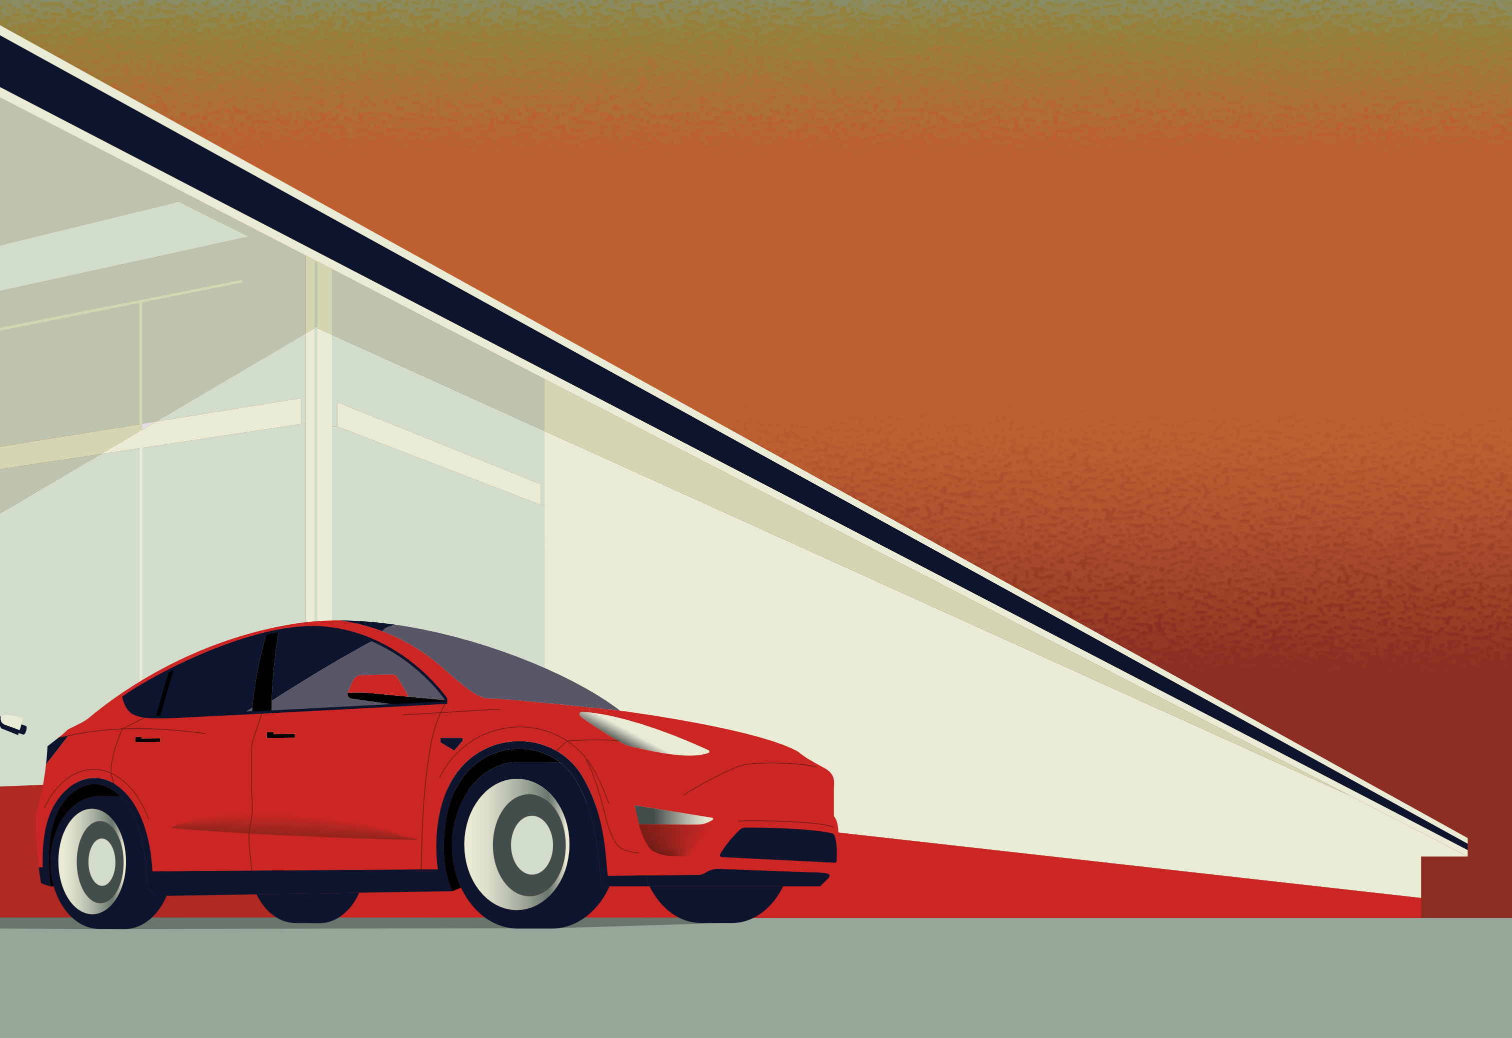

ON THE A13 HIGHWAY IN GERMANY, a harried mom pulls into a rest stop, her two restless teenagers in the back of the family's electric SUV. She steers toward a row of 24 sleek, refrigerator-size obelisks, most already tethered to a vehicle, and parks in front of the unoccupied one she'd reserved en route. Unhooking the cable, she inserts the plug on its end into a port in the car's rear left flank, behind a flap that resembles the fuel door of earlier decades. She and the teens head to the bathrooms and a warm café for the 15 minutes it'll take to recharge the car.

On the way, they glance over at the travel plaza's fossil-fuel section. It's like a little slum, with oil stains on the tarmac, the smell of petrol fumes in the air, and drivers standing at their vehicles squeezing the grimy handles of fuel nozzles. At the fast charger she's using, on the other hand, the whole process requires no human intervention beyond plugging in. Her SUV identified itself to the charging kiosk, her charging network authorized payment, and voilà, the torrent of electrons began. The cost of the recharge will be added to the amount the charging network deducts monthly from her payment account.

The 15-minute recharge gives the SUV another 375 kilometers (235 miles) of range; it might have been a few kilometers more if the day weren't quite so chilly. Still, it's more range than our driver will need to get to Prague, where she and the kids will meet her husband, already in the city on business.

This is the fantasy of the electric-vehicle future a decade from now. Will it come to pass? Hard to say. Fast charging is evolving quickly, and the technology, standards, and business factors that will shape the outcome are only now coming into focus. Certainly, there's still a lot to be done. And the stakes are high: A McKinsey & Company report suggests that over the next decade, more than US $50 billion will be spent on EV charging sites.

Porsche's first fast-charging station, at the Berlin Adlershof Technology Park, recently went live.

Tesla was the first to get a sizable network of fast chargers up and running. The company now operates some 1,500 stations with more than 13,000 charging cables in North America, Europe, and China. Tesla unveiled its first 90-kilowatt Supercharger station in 2012, showing the world what a modern, pleasant, fast-charging site should look like. By 2015, Tesla drivers could cross the continental United States solely on electricity, with their cars' navigation systems routing them through a succession of Supercharger sites, even telling them how long to stay plugged in to minimize wait time.

The rest of the industry intends to catch up in time for a new wave of long-range electric cars to hit the market in the early 2020s. Major car companies have aligned themselves with three competing fast-charging standards. The standards they're promulgating are incompatible, each having its own operating voltages and currents, physical connectors, control logic, and communication protocols between car and charger.

Nissan and Mitsubishi are behind a Japanese system called CHAdeMO, for "charge de move," which is also an allusion, in Japanese, to the phrase "a cup of tea," referring to the time required to recharge the vehicle. China, maker of two-thirds of the world's lithium-ion batteries, has developed a standard called GB/T, for *Guóbiāo tuījiàn* (from the romanization of the Chinese term for "preferred national standard"). Meanwhile, all European and U.S. manufacturers except Tesla have lined up behind a standard called Combined Charging System (CCS), and even Tesla has said its Model 3 vehicles sold in Europe will have a CCS port.

The CCS group includes BMW, Daimler, Ford, GM, Hyundai, Jaguar, Kia, Renault, and, notably, VW Group. And within the VW Group, Porsche, the high-performance brand that's arguably the group's crown jewel, has made little secret of its intention to do Tesla one better.

PORSCHE

DUELING CHARGERS

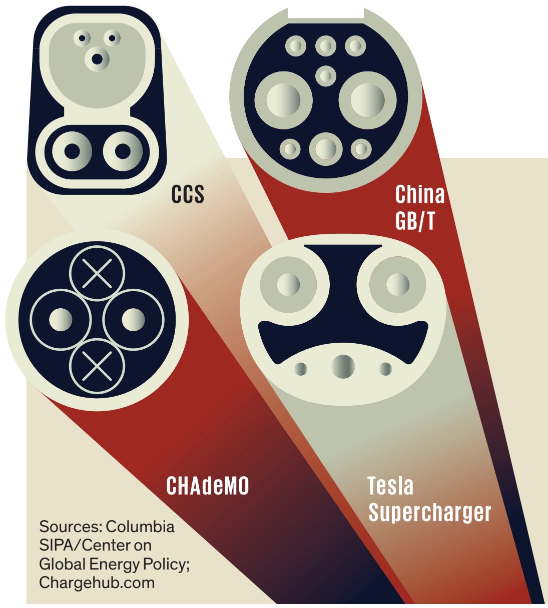

Sources: Columbia SIPA/Center on Global Energy Policy; Chargehub.com

System	Public chargers worldwide	Kilowatts	Availability
Combined Charging System (CCS)*	22,000	50–350	United States, European Union, Australia, Korea
China GB/T	330,000	237.5	China, India
Tesla Supercharger	13,000	135	Global
CHAdeMO	25,300	50–100	Global

* North American and European versions are not compatible.

Porsche's new Taycan, the brand's first all-electric car, was revealed to the world in production form as this article was going to press. The price will start around $90,000 and go as high as $130,000, close to the Tesla Model S pricing these days. The sedan's performance roughly equals that of Tesla's Model S. However, it will offer fast charging that's much faster. Tesla's Supercharger sites now operate at up to 135 kW, meaning a Tesla driver can add up to 320 km (200 miles) in as few as 30 minutes. The company plans to boost peak power at its sites to 250 kW over the next few years.

Final details will be made public only on 4 September. But if published reports are accurate, the Taycan will launch with 250-kW charging, and 350 kW will be possible by 2021 at the latest. That means 400-plus kilometers of range can be added in less than 20 minutes. In a real hurry? Figure on 100 km in 6 minutes.

Those levels of power demand the use of new hardware never before fitted to a passenger car. Besides more powerful charging stations, the system also requires a grid that can carry higher power levels to the charging sites, as well as many other enhancements. To roll all this out on a large scale, dozens of car companies will have to work with hundreds of charging networks, equipment suppliers, and, indirectly, with thousands of electric utilities around the world.

US $50 billion will be spent on EV charging networks in the next decade.

TODAY'S EV CHARGING isn't all that fast. It's certainly not as quick at adding range as pumping gasoline. But that doesn't matter so much, because most recharging is done overnight at home or during the day at the workplace. In North America, battery-electric vehicles are usually charged at 240 volts AC, using what are known as Level 2 charging stations. (Level 1 is the region's standard 120-V wall current, adequate for adding 30 or 40 km overnight.) Depending on the car's built-in charger—Tesla aside, they mostly range today from 3.3 to 7.2 kW—it can add 15 to 50 km (10 to 30 miles) each hour. In one night that's enough to fully recharge a typical 60-kilowatt-hour EV battery pack, if it's not completely depleted at the start.

Fast charging is a different story. First of all, it uses direct current instead of alternating. Other than Tesla, the electric cars on the road today are mostly limited to 50 kW, meaning it takes most of an hour to charge that nearly tapped-out 60-kWh battery to 80 percent capacity or up to 90 minutes for the larger 90-kWh battery packs in luxury vehicles.

The changes in the charging rate as the battery charges also affects the time a driver spends plugged in. Some cars, among them the Chevrolet Bolt EV, start to taper down the charge rate once the battery's state of charge hits 50 percent.

The 125-kW Tesla Supercharger system also tapers, but today it provides up to 300 km in 25 to 45 minutes.

Those differences barely hint at the lack of uniformity in today's charging networks. North America has a dozen or more individual public charging networks offering a mix of Level 2 and fast charging. Generally, drivers must belong to a network to use its stations. If not, they have to call a toll-free number and provide a credit card to an agent, who unlocks the station. Virtually no stations have credit card readers, as gasoline pumps do. Experienced EV drivers may have tags, fobs, or apps from as many as eight separate networks. The kind of roaming that mobile-phone operators enabled 15 years ago doesn't exist in the United States, although Canada and Europe are considerably better in that regard. In 2019, several U.S. networks announced plans to allow future roaming with other individual networks—but those are bilateral agreements. The industry is still well short of cellular-style transparent nationwide roaming.

When Porsche set out to define its first electric car, it surveyed its customers about what they would want in a fast, pricey, high-performing electric sport sedan. As well as consistent performance throughout the battery range, the company found that buyers wanted recharging on road trips that didn't take much longer than filling up with petrol. Porsche set a 15-minute target, 20 minutes at the outside. Payment had to be at least as easy as using a gasoline pump at a highway rest stop.

That's what Porsche says it will deliver, sometime in 2021, in its Taycan and the 350-kW charging stations to follow that only the Taycan will be able to use. That exclusivity will be fleeting; other pricey EVs will incorporate the same standard. But Porsche will have achieved its goal of being first (and of beating Tesla).

Doubling the operating voltage of the Taycan, to 800 volts, should save about 30 kilograms.

Tesla, on the other hand, asserts that such extreme fast charging isn't required. A company executive, interviewed on the condition that he not be named, said that surveys suggest Tesla owners value the quality of their recharging experience more than its brevity. Customers might relax in a café, go shopping, or simply chat with other Tesla owners. However, the company is starting to upgrade its existing Supercharger network to provide an average of 175 kW per car, or about half of what Porsche is promising, with a peak rate as high as 250 kW per car.

To learn more about how Porsche accomplished this remarkable piece of engineering, I visited the heart of its research operation this past March.

THE TWO-LANE ROADS between the outskirts of Stuttgart and the hamlet of Weissach rise and fall through sweeping, hilly curves with open fields on each side. They have no shoulders, which heightens the driver's senses while moving in traffic that runs at precisely the marked limit of 100 km/h (62 mph). But the pace regularly slows as heavily loaded tractor trailers and car carriers packed with new Porsches navigate the bends–suggesting that Weissach isn't your typical German farming town.

Indeed, it's been the home of the Porsche R&D center for nearly 60 years. These days, it's a hive of construction, with barricades, cranes, and road diversions that lead to new and expanded parking lots terraced into the hillside. From behind a tall wooden fence blocked by shrubbery and posted with strict warning signs comes the howl of cars tearing along the company's closed high-speed test track.

I'm here on this sunny spring day to speak with Joachim Kramer, who leads Porsche's work in high-voltage electronics for electrified vehicles. With the official Taycan launch just five months away, he's a very busy man. His answers to many questions are guarded, to conceal specifications still being confirmed and also to ensure he doesn't reveal to competitors any of the lessons Porsche has learned in the seven years since it first got a good look at a Tesla Model S.

Kramer is confident Porsche will deliver the capabilities it has promised. The engineers encountered "no real surprises," he says–though a few unexpected aspects required some novel engineering. The engineers faced a trade-off, he says, between battery capacity and charging speed. The bigger the battery, the less often charging is needed, but the pricier the car. They concluded the largest batteries in the most expensive electric cars will remain at 90 to 100 kWh for the next several years. "In 10 years, perhaps 200 kWh?" he muses.

That meant the power had to go up significantly. The company's previous hybrids and plug-in hybrids operated at 300 to 400 V, the standard industry prac-

MARY TURNER/REUTERS

CHARGING GIBBERISH

Ask anyone what kind of fuel their car takes, and you'll likely hear "gasoline." Ask for more details, and the most you'll get will be "regular" or "premium." Now ask what kind of charging an electric car might have, and you'll get a blank stare.

Explaining charging to car shoppers is hard. It sounds to most car buyers like you're reading from an engineering specification. In fact, most of the following terms come from just such a list of specs. They are, nonetheless, the bare minimum that today's modern electric-car buyer must know to sort out where a car can and cannot charge.

J1772: SAE International's protocol and connector standard for Level 1 and Level 2 charging. Used on all electric cars except Teslas. (Varies between North America and Europe.)

Level 1: 120-volt AC charging in North America (the residential standard), via a charging cord with a J1772 connector.

Level 2: 240-volt AC charging in North America, also J1772.

CHAdeMO: DC fast-charging standard, used only for Japanese cars and a few early Korean models. Tesla offers a CHAdeMO connector for its cars.

CCS, or Combined Charging System: DC fast-charging standard, used by all European makers and all U.S. makers except Tesla. Its global use is promoted by a group called CharIN.

Supercharger: DC fast-charging standard usable only by Tesla vehicles.

No automaker, charging network, or standards body has proposed any way to simplify this gibberish into an easily understood, graphically intuitive system for signage, mapping, and vehicles. The closest anyone has come is the Chargeway graphic system privately developed by marketer and graphic designer Matt Teske. It uses a series of shapes, numbers, and colors to distinguish among types and speeds of charging. It is now being tested by the Oregon Department of Transportation at some pilot locations.

tice. But if you want your fast charger to use that voltage you'd have to boost the current substantially. That would mean increasing the size of the copper wires, adding more mass.

So the company doubled the voltage by connecting pairs of cells in series. With this 800-V system, Kramer estimates Porsche saved roughly 30 kilograms in the Taycan—a worthwhile amount even in a vehicle likely north of 2,000 kg. That mass savings is divided among the battery pack, power electronics, windings in the permanent-magnet motors that power the wheels, and the thick cables that connect it all together.

At a component level, the higher voltage required extra insulation, including wider air gaps between components, as well as other features. However, none of Porsche's traditional suppliers offered such components. Those from other suppliers—MOSFETs for electric locomotives and other uses, for instance—weren't qualified for the environmental conditions, temperature swings, vibration tolerance, and 10-to-15-year life required for road vehicles.

The Taycan will identify itself to the charger, so no other form of payment will be needed.

THIS IS NO ORDINARY SUPPLY PROBLEM. Very few, if any, high-volume mass-market products handle these levels of power and voltage. Magnetron tubes in microwave ovens, for instance, require a power supply of 1,000 V or more, but only at 1 or 2 amperes.

So Porsche had to work closely with suppliers to prototype, test, develop, and retest scores of new components. Kramer gets vague when asked about specific issues its engineers faced, but he identifies two in general terms: energy storage and electromagnetic radiation.

Both are strictly limited by vehicle regulators around the world. By law, an EV must be able to drain energy quickly from its high-voltage system, because in a crash that energy could shock and injure emergency responders cutting into the car to rescue occupants. Kramer says limiting energy storage requires diverting energy into multiple components—especially capacitors—that can be shut down simultaneously and rapidly.

As for radiation, Kramer will say only that the company identified it as a challenge in its higher-voltage prototypes, and that the design has shielded against it appropriately. Beyond preventing radiation from interfering with the vehicle's electronics, regulators worldwide limit the radiation to levels they consider safe for human occupants and those in the vicinity of the vehicle. Kramer says simply that Porsches meet all such global regulations.

Kramer confirmed that Porsche will use the CCS standard in all markets except Japan, where the CHAdeMO standard is transitioning toward higher power levels, and in China, where Porsche Taycans will adhere to the GB/T standard. (Irritatingly, the sockets and connectors for the European and North American versions of CCS differ just enough that one region's version can't charge cars built to use the other.)

25 percent more electric cars are sold in China than in the rest of the world combined.

The CCS 2.0 specification, a superset of the SAE J1772 protocol and connector in North America, is designed to transfer up to 1,000 V into the vehicle—depending, of course, on the voltage its particular battery can accept. Porsche's early tests suggested the weak point in existing EV hardware would be the charging sockets, which might not reliably tolerate such high voltages. So new components are required there, too.

Any Porsche model undergoes a sequence of performance improvements over its lifetime to keep customers engaged. The Taycan will be no different, Kramer says, hinting that more power from the battery to the motors is not only possible but likely.

"We can go higher," he says, smiling. The challenge? "It becomes hard to drive" because it has so much power. Consider the Tesla Model S P100D, which is already capable of reaching 100 km/h in 2.3 seconds from a standing start. Even die-hard Porsche enthusiasts accustomed to very high levels of performance may find there's a limit to their abilities to handle the power they crave—and to their passengers' tolerance of the resulting *g*-forces.

EV Car Predictions for 2040

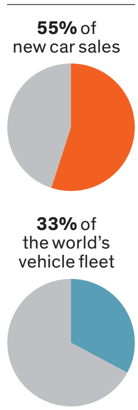

Source: Bloomberg New Energy Finance 2018 Electric Vehicle Outlook

EVEN LUDICROUS PERFORMANCE won't sell an EV if buyers can't conveniently recharge wherever they may want to go. Two years ago, a group of automakers in Europe (BMW, Daimler, Ford, VW Group) recognized the problem and set out to solve it.

In Western Europe, the result is the network of high-speed charging stations spreading fast under the Ionity brand. It was established in November 2017 as a follow-on to CharIN, a group of automakers and equipment makers formed to develop and promote use of the CCS fast-charging standard. Since 2018, the Ionity network has opened over 100 fast-charging stations toward its goal of 400 by 2020. At these stations, an electric-car driver can stop for 15 to 45 minutes and recharge a vehicle as fast as its onboard charger permits—all the way up to 350 kW.

The Porsche Taycan will also be the first vehicle launched with a feature called Plug and Charge (ISO 15118), a standard under which the vehicle identifies itself to the charging network so the driver doesn't need to present a membership card or payment method. Instead, on the back end—with personally identifiable information shielded by layers of public-key infrastructure encryption—the network being used identifies the vehicle, its owner, and the payment method. Then it creates a transaction that's seamlessly charged to the specified account.

The standard has to be incorporated into both the vehicle and the charging site, with back-end modifications to permit seamless roaming. The Ionity network in Europe will have it all integrated in time for the arrival of the first cars that can use it.

In North America, Electrify America (EA) is the only network so far that has announced it will include Plug and Charge. In May 2019, that network's technical center in Herndon, Va., held a test day to give carmakers and charging-station manufacturers the chance to test their hardware via EA's network to ensure compatibility. Similar efforts from other networks will follow over time. Indeed, North American networks have issued bilateral roaming and interconnection agreements regularly through the first half of 2019.

The roots of Electrify America date to September 2015, when the U.S. Environmental Protection Agency announced that VW Group had installed software in diesel vehicles that disabled emission controls on the road—that is, outside of EPA lab tests. The scandal, which involved about half a million vehicles sold in the United States from 2009 through 2015, overturned VW's long-standing bet on diesel. That's what led the company to embrace what it had only dabbled in: electric drive.

VW Group has so far allocated €30 billion to cover the costs of the diesel emissions scandal. The impact is global, also affecting 11 million diesel vehicles the company had sold in Europe. In a comprehensive settlement of multiple charges by various U.S. and state agencies, VW provided about $2 billion to establish a fund to provide zero-emission vehicle infrastructure in the United States, with 40 percent allocated to California. The funds were to be spent in four 30-month segments, with approvals by the EPA and the powerful California Air Resources Board required at each interval.

The result is the Electrify America charging network. The agreement states that by 1 July 2019, the network was to have more than 2,000 fast-charging cables in 484 discrete sites operating or under construction throughout the United States. The pace was rapid; its first station, in Chicopee, Mass., didn't open until May 2018. Nevertheless, EA met the goal. (There's an Electrify Canada too, separately funded.)

Though the EA network is a separate corporate entity, it is wholly owned by the Volkswagen Group of America, tying the company to its wholesale transition into battery-electric vehicles. It says it plans to build and sell 1 million of them a year by 2025, out of total global production of 10 million vehicles. Volkswagen asserts such plans were under way before the EPA scandal, but observers are skeptical.

That million-EVs-a-year figure can be deceptive for Americans, however. The bulk of those vehicles are likely destined for China, and to Europe thereafter. While VW plans to assemble a battery-electric compact crossover utility vehicle in Tennessee, starting in 2022, analysts suggest that North America will lag significantly in its adoption of battery- electric vehicles.

ChargePoint stations (projected growth)

Source: ChargePoint

THERE ARE MANY REASONS FOR THE LAG. One is that unlike Europe or China, the United States has more than 3,000 separate electric utilities, overseen by 50 state public-utility commissions. California, for instance, has more than 50 utilities, split among larger investor-owned companies, publicly owned providers, and rural electric cooperatives. That makes it hard to scale up the installation of charging infrastructure. By contrast, in many European countries Ionity was able to negotiate with only a single national electricity provider.

Another reason is that U.S. drivers drive farther, on average, than those in Europe, Japan, or China. Gasoline in the United States is cheap by global standards, and there is a dearth of mass transit between cities 300 to 500 km apart, a distance that is usually too short to fly economically.

Finally, North American automakers have yet to persuade the industry they're fully on board with battery-electric vehicles. General Motors is by far the furthest along. Its CEO, Mary Barra, said in 2018 that GM expected to be able to sell electric vehicles with ranges of 320 km or more at a profit during the early 2020s. The company's last new electric vehicle launched in December 2016, however, and its next all-new one isn't expected until after 2020.

As for the others, Ford hasn't introduced a new electric or plug-in hybrid vehicle since 2013. It says it will have a 300-mile (482 km) electric crossover utility for 2020. And Fiat Chrysler, perennially seeking a merger partner among the world's automakers, has bigger challenges than electric cars.

U.S. automakers have declined to follow Tesla's lead in setting up charging networks. Throughout its first several years of selling the Volt plug-in hybrid and Bolt EV electric car, GM said it had no intention of spending any money to provide charging infrastructure. A recent announcement that it would join Bechtel Corp. in a joint venture to do just that represents a change in tone, but the project's scope and impact remain to be seen.

THE BIGGEST FACTOR in the world's electric-car race, however, is China. Its government has long wanted to dominate global supplies of lithium-ion battery cells, photovoltaic solar cells, and plug-in electric vehicles. In 2018, 1.2 million electric cars were sold in China, about 25 percent more than in the rest of the world combined. Total global vehicle sales number roughly 86 million.

GB/T, China's own fast-charging standard, is currently specced at more than 200 kW and up to 1,000 V. In August 2018, the China Electricity Council agreed to partner with the Japanese CHAdeMO consortium for a standard, now known as ChaoJi, that can handle 900 kW.

But the country's fast-charging infrastructure today is less developed than that in Europe and Japan, says longtime China hand and auto-industry consultant Michael Dunne of ZoZoGo.

"Chinese consumers are superpractical and frugal," Dunne explains. "Many are happy to charge during the day at the office. Standard charging and free, or lower cost, is much preferred over fast charging but expensive." He cites Tesla as the country's current leader in fast-charging sites, with more than 250 Supercharger locations in China at the start of 2019.

Still, China is known for doing things at scale. It now has what is purported to be the world's largest charging site, the Minle site in Shenzhen. In May, Southern Power Grid added another 172 fast-charging cables to bring the total to 637 fast-charging connectors that can charge up to 5,000 vehicles a day.

Tesla, Dunne notes, also understood that even wealthy Chinese buyers often live in apartment buildings. It set up urban Supercharger sites to allow them to recharge once or twice a week even if they don't have a dedicated parking space where they live or work.

Porsche's Kramer echoes that theme as well, suggesting that by 2025 or so, 350-kW fast charging will spread to less expensive and higher volume VW Group vehicles. That charging rate would similarly make electric cars practical for apartment dwellers in European cities even if they park at the curb, in public parking structures, or in open lots.

It tends to be more complex, time-consuming, and expensive to build high-voltage charging stations in urban centers than in the countryside. It would seem that the companies that are best placed to do it are the electric utilities that already serve those areas.

What role utilities might play in charging cars could be the subject of several articles. However, despite uninformed suggestions that electric-car charging demands may crash the grid, the utilities of the world are unfazed by the prospect of EV charging plazas drawing up to 2 megawatts from a few dozen stations.

A standard-issue big-box retail store—Walmart, say—requires up to 2 MW for things like refrigeration and lighting. No one raises the specter of thousands of Walmarts crashing the grid. Charging stations for EVs are designed and approved like anything else of size that gets built, and electric utilities know how to provide the required electric capacity to new construction. It's what they do.

The role of utilities in building and operating charging sites themselves, however, varies considerably. For years, utilities in most areas outside early-adopter California watched EVs cautiously. Now, major U.S. utilities are recognizing that the market will come and that all types of charging offer a rare chance to increase their operations and footprint. Moreover, in many cases, their regulators may let them pass the costs of such expansion to consumers, a practice known as rate-basing, if they can demonstrate a public benefit.

Unsurprisingly, privately funded charging networks cry foul. A number of lawsuits have been filed over the past five years to block or modify submissions by utilities to set up charging networks in their service areas. By and large, especially in California, such operations are being approved—with strong provisos requiring installations in underserved communities that commercial operators may deem unprofitable.

WHEN PORSCHE TAYCANS HIT THE ROAD—late this year in Europe, early in 2020 in North America—they'll be bought by affluent people, who generally own their own homes. That's where they'll charge the cars, for the most part. On long road trips, Electrify America and soon other providers—EVgo and ChargePoint are two—will provide very fast charging as needed, and their networks will expand in coming years to meet demand.

The question is whether there's a business there. Or might EV charging be a necessary service that must be provided somehow to sell cars and cut carbon emissions? Electric-car advocates concluded years ago that there's no business model in selling electricity via public 240-volt Level 2 charging stations. These seem likely to become ubiquitous over time. A good parallel might be Wi-Fi service at hotels and public spaces like airports: It's often free, though providers will charge for it where they can.

The jury is still out on whether there's a business model for DC fast charging. A rough benchmark seems to be that it shouldn't cost much more per mile than gasoline. As analyst John Gartner of Navigant Research notes about the added costs of providing more electric supply to a site, "If it's challenging to get cost recovery at 25 and 50 kW, what kind of utilization do you need to cover your costs at 150 or 350 kW?"

Still, these are early days, and lots of companies are trying lots of approaches. And consider the changes over the nine years that modern electric cars have been on sale.

Business analysts either scoffed or scratched their heads when Tesla first announced that it would build its own dedicated high-speed charging network in the United States and globally. But it did exactly that, and in just six years. As in so many other aspects of electric cars over the last decade, Tesla turns out to have been prescient.

The history of technology provides ample evidence that those who pioneer technology often don't profit from it—or even survive. Whether or not Tesla remains an independent carmaker 10 years hence, it showed what was possible. The rest of the world is now following in its footsteps. ■

↗ POST YOUR COMMENTS at https://spectrum.ieee.org/fastcharging0919

Countering Counterfeit Drugs

By Swarup Bhunia & Soumyajit Mandal

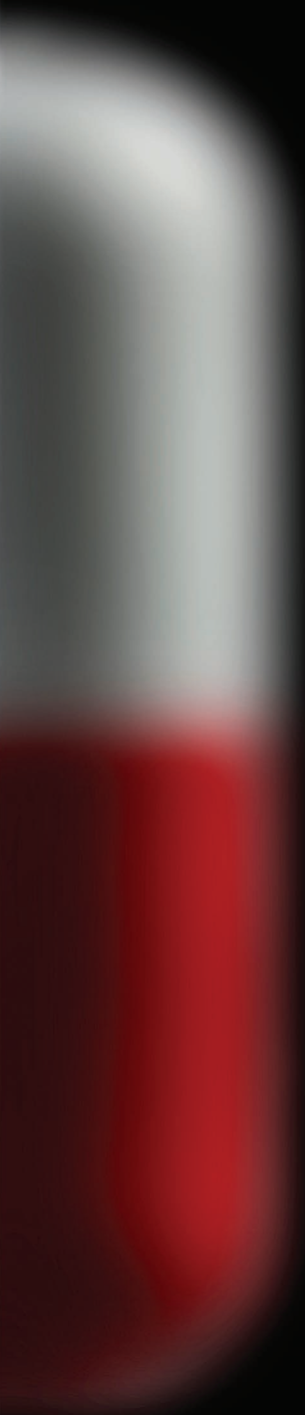

ADAM VOORHES/GALLERY STOCK

A technique used for detecting explosives can also verify the integrity of medicines

When you purchase medicine at the drugstore, you assume that it's what you think it is and that the active ingredient in the drug is present in the specified concentration. Unfortunately, your assumption might be all wrong. Counterfeit and substandard medicines have become widespread, particularly in low- and middle-income countries with weak regulatory systems. Indeed, according to the World Health Organization (WHO), one out of 10 medicines sold in developing countries should be considered "substandard." Your drug could even be an outright fake.

"But I live in the United States," you may say. "The medicines at my pharmacy are regulated by the U.S. Food and Drug Administration, so it must be the genuine article." Unfortunately, even the United States and other higher-income countries aren't immune to this scourge. Since 2012, smugglers have been caught selling fake drugs to more than 3,000 doctors, clinics, and hospitals across the United States.

In one notorious case, two lots of the cancer drug Avastin were discovered to contain none of that medicine's active ingredient. After a recall and lengthy investigation, the FDA concluded that the fake Avastin had traveled through a network of overseas suppliers, passing through Canada before reaching the United States.

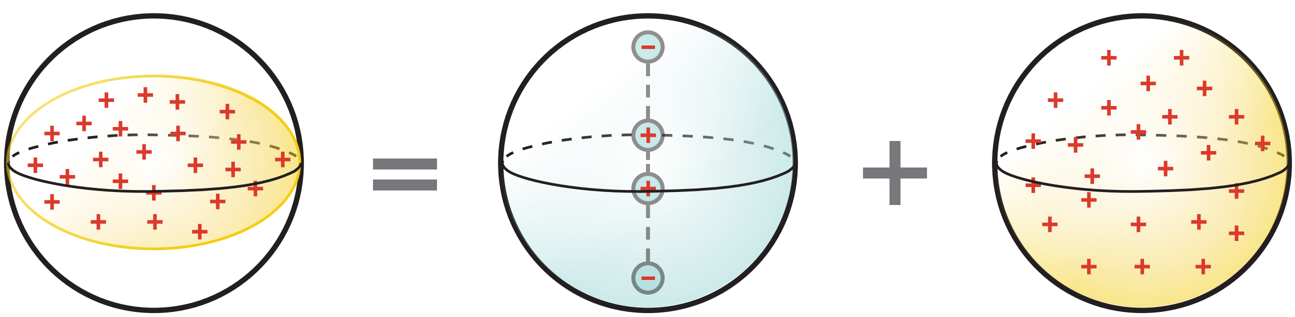

FOUR POLES: Nuclear quadrupole resonance requires an atomic nucleus with a nonspherical distribution of positive electric charge [left], which creates an electric quadrupole moment. The word *quadrupole* refers to the four electric poles [center] that produce an equivalent nonspherical charge distribution when added to a set of spherically distributed charges [right].

And even if your drugs aren't counterfeit, what assurance do you have that the pharmaceutical ingredients have not degraded? Many drugs are sensitive to heat, and neither you nor your pharmacist has any way of telling whether the pills you've just picked up experienced problematically high temperatures—say, in the back of the truck that delivered them.

The growing use of online pharmacies is only making these problems worse. In the United States, tens of millions of people (about 8 percent of the adult population) buy medicines outside the formal drug supply chain, typically from foreign online pharmacies or other unlicensed sources. A quarter of the people in the United Kingdom say they're likely to use an online pharmacy in the near future.

The issue of fake and substandard medicines has attracted the attention of law enforcement authorities around the world. In 2008, Interpol created Operation Pangea, a special division targeting the online sale of illegal pharmaceuticals; last year it seized 10 million units of fake and illicit medicines and shut down more than 3,600 websites. And hundreds of millions of doses are seized every year at international borders. The global market in counterfeit drugs is estimated to be worth somewhere between US $75 billion and $200 billion annually. People are wasting a lot of money buying this useless stuff.

But forget the economic losses: The human toll is what really matters, and it is enormous. Ineffective antimicrobial drugs in particular compromise the treatment of many deadly diseases, especially in poor countries. They also promote development of antimicrobial resistance. WHO estimates that 72,000 to 169,000 children die each year from pneumonia due to substandard or counterfeit antibiotics. And in sub-Saharan Africa, an estimated 64,000 to 158,000 people die every year from malaria that they tried to treat with substandard or fake antimalarial medications.

ILLUSTRATION BY Erik Vrielink

It's no exaggeration to compare illicit medicines with the nuclear and biological weapons of mass destruction we all fear. These WMDs, though, are largely aimed at people in poor countries who are already facing a multitude of social and economic ills.

Clearly, we need better ways to detect medicines that are fake, substandard, or degraded from negligent handling. Efforts to improve the situation include recent provisions of the European Union's Falsified Medicines Directive, adopted last year. As of last February, all packaged medicines transported into and across the EU must now have a "unique medicines identifier"—a separate reference number for every packet of medicine, not just every brand or batch—and also a tamper-evident seal. In the United States, the Drug Supply Chain Security Act, signed into law in 2013, requires all packages of medicine to be electronically traceable by 2023.

Sadly, these "track and trace" approaches suffer from a fundamental shortcoming: They authenticate the package, not its contents. A package of medicine is assumed to be genuine simply because it has a valid security mark. Unscrupulous manufacturers can readily circumvent such measures by putting the wrong stuff in the right package. And these approaches offer no help in detecting degradation.

Our research groups at Case Western Reserve University, in Cleveland, and the University of Florida, in Gainesville, along with some colleagues at King's College London, have been looking at a technology that promises to be a whole lot more effective—offering a way to authenticate the actual contents of a packet of medicine. You might imagine that any system used to verify the chemical constitution of a dose of medicine would necessarily destroy it. In fact, though, using a physical phenomenon called nuclear quadrupole resonance (NQR), you can test your pill and eat it, too.

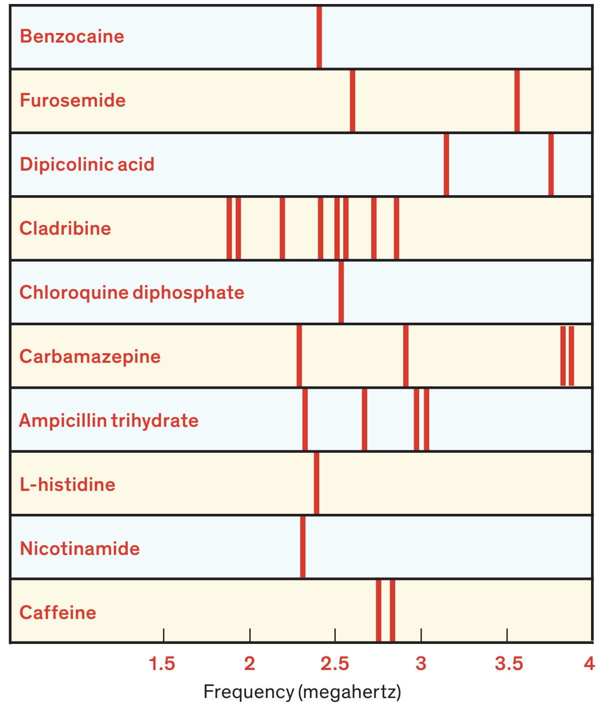

ON THE RADIO: Compounds that contain nitrogen, including the drugs shown here, produce distinct NQR spectra with radio-frequency resonances [vertical bars] in the range of a few megahertz. Chlorine-containing compounds produce higher-frequency resonances.

NQR provides a relatively easy and nondestructive way to perform a chemical analysis on many kinds of packaged medicines and dietary supplements. The electronic tester we've been developing measures the NQR response of the sample at different frequencies. The resulting NQR spectrum is generated by energy transitions within the atomic nuclei of the chemical, providing a unique fingerprint for that compound. While the quantum mechanics of the phenomenon is complicated, here's what's basically going on at the atomic level.

Atomic nuclei contain positively charged particles (protons), usually along with electrically neutral particles (neutrons). If the distribution of positive charges is not perfectly symmetrical, the nucleus will have what is called a charge moment, which characterizes how the charge is distributed. In the nuclei of some atoms, the positive charge, rather than being spherically distributed, stretches toward two poles of the nucleus, creating what is known as an electric quadrupole moment. (It's not possible for a nucleus to have an electric dipole moment.) Or the positive charge could accumulate around the "equator" of the nucleus, which would also give rise to a quadrupole moment (but of the opposite sign).

Nuclei with quadrupole moments can occupy only certain distinct energy levels. Those levels are determined by the interaction between a nucleus's quadrupole moment and the charge distribution of the electrons that surround the nucleus. That charge distribution is in turn determined by the chemical environment the nucleus finds itself in. So by measuring the energy levels in the nucleus—or more accurately, by measuring the differences between energy levels as the nucleus shifts from one level to another—you can infer what the chemical environment is.

About half of the elements in the periodic table have some isotopes with nuclear quadrupole moments. Many of those isotopes are rare. Fortunately, the most common isotopes of nitrogen (nitrogen-14) and chlorine (chlorine-35 and chlorine-37) have nuclear quadrupole moments, and many medicines contain one or both of these kinds of atoms.

When subjected to radio waves of between 0.1 and 5 megahertz, nitrogen nuclei will shift energy levels,

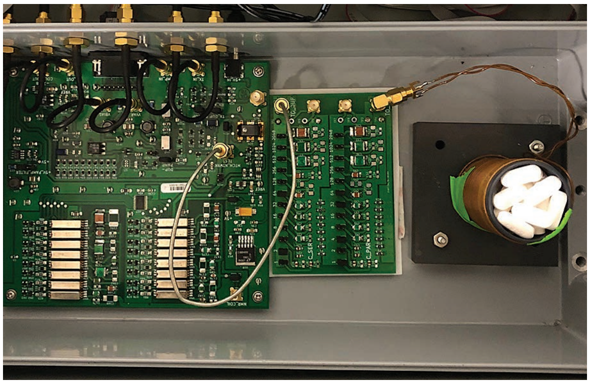

INNER WORKS: The authors' prototype instrument for measuring nuclear quadrupole resonance contains a simple coil [right, with pills inside] and two analog signal-processing boards [left and center]. Digital signal processing is done by circuitry located in a second box.

absorbing and later reemitting radio energy at distinct frequencies that depend on the chemical and, to some extent, the physical environment of those nuclei. Chlorine nuclei do the same thing, but at frequencies from 20 to 40 MHz.

NQR spectra in these frequency bands can thus serve as unique chemical fingerprints for nitrogen- and chlorine-containing compounds. By measuring such spectra, you can readily identify many kinds of drugs. Indeed, in some cases NQR spectra can even be used to distinguish between identical drugs prepared in slightly different ways–say, tablets containing different inactive ingredients or produced using machines that apply different compacting pressures. Our NQR-based drug-testing apparatus has, for example, been able to reliably classify tablets of acetaminophen, an over-the-counter pain killer, made by different manufacturers.

Nuclear quadrupole resonance is useful for testing specimens that are solids or powders, but not liquids. While that's an obvious limitation, NQR has a lot of other things going for it. In particular, it's insensitive to the presence of coatings or packaging materials. So it can be used to examine pills while they're still in the bottle or blister pack. Indeed, it could be used to test an entire shipping carton of such bottles or packs, or a drum of powdered material.

What's more, the equipment could be built at low cost and would be amenable to miniaturization. And because NQR instrumentation relies on radio waves of relatively low frequency and power, it is inherently safe and could be used without special training.

You may have noticed the similarity between the terms *nuclear quadrupole resonance* and the more familiar *nuclear magnetic resonance*, which is the physical phenomenon on which magnetic resonance imaging (MRI) is based. Indeed, there are many similarities between the two phenomena: Both are sensitive to the presence of certain kinds of atomic nuclei, both detect transitions between nuclear energy levels by measuring radio signals in the megahertz range, and both can be used to produce images.

A key difference, though, is that NQR doesn't require that the sample be placed in a strong magnetic field. So there's no need for an expensive superconducting magnet of the type found in MRI scanners. All you really need is a coil of wire and some suitable electronics for generating the appropriate radio-frequency excitations and measuring the sample's response.

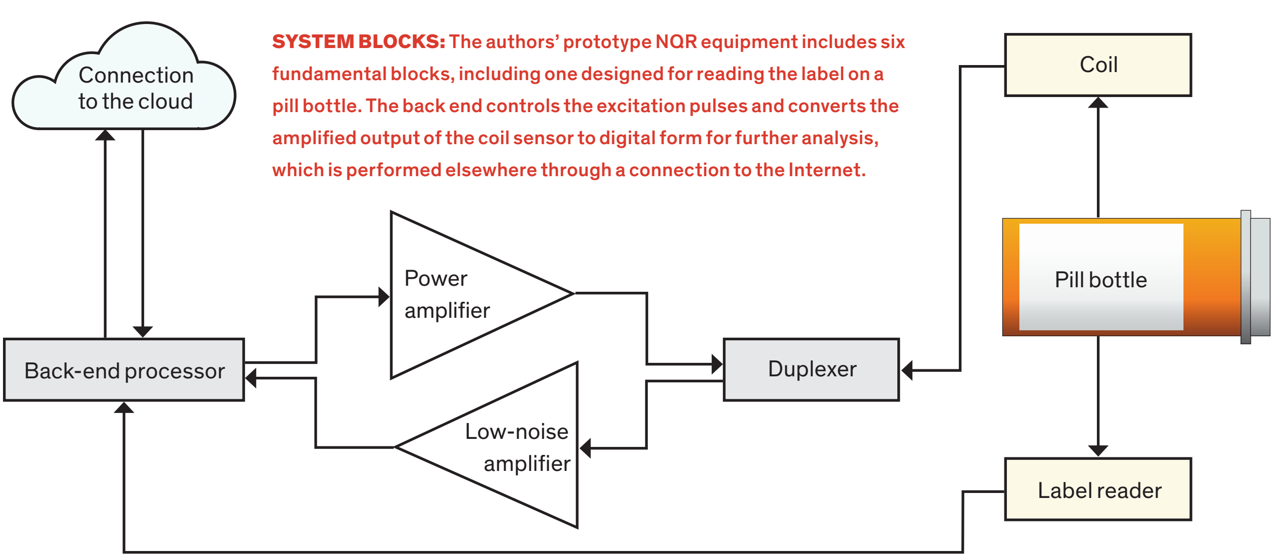

SYSTEM BLOCKS: The authors' prototype NQR equipment includes six fundamental blocks, including one designed for reading the label on a pill bottle. The back end controls the excitation pulses and converts the amplified output of the coil sensor to digital form for further analysis, which is performed elsewhere through a connection to the Internet.

SWARUP BHUNIA AND SOUMYAJIT MANDAL

The excitation consists of one or more RF pulses generated by a power amplifier producing at most a few watts. These pulses are applied to the sample by a transmitter coil that generates an oscillating magnetic field. The field gives rise to oscillations within some of the atomic nuclei in the sample; the oscillations can be measured by a receiving coil attached to a suitable RF amplifier. As a result, the equipment is pretty simple. And you can simplify it further by including a switch that allows you to use the same physical coil for both transmission and reception.

The electronics do have to be sophisticated enough to amplify and process the acquired NQR signal and extract parameters of interest: signal amplitude, frequency, phase, and decay rate. In the instrument we have built, an embedded computer compares these parameters against reference data to classify the sample. Our prototype drug-authentication device is portable, performs measurements automatically, and doesn't require any special skills to operate. It could thus be used anywhere in the drug supply chain. We estimate such a device could be manufactured at a cost of about $100, which would, presumably, translate to a price for the end user of less than $1,000.

Our prototype is portable, performs measurements automatically, and doesn't require special skills to operate.

In our view, NQR holds much more promise than any other analytical technique you could imagine. Liquid chromatography and mass spectrometry, for example, while effective at identifying chemical compounds, involve expensive apparatus and necessitate destroying the sample. Optical methods, such as near-infrared or Raman spectroscopy, don't have those drawbacks, but they're able to probe only the surface of a drug sample, so they couldn't be used for a pill that's still in its packaging or for medicine inside a capsule. And these optical methods could be fooled by pills manufactured with coatings that contain the features of the correct drug.

That said, we readily admit that NQR has some special challenges. The biggest are its inherently low signal levels and poor signal-to-noise ratios. Those same issues are what stymied earlier applications of NQR to explosives detection. A related problem is external RF interference, created, say, by an AM radio station or by a noisy switching power supply in nearby equipment. And finally, the NQR sample coil is normally attached to capacitors in what's known as a tuned circuit, so that power can be applied to it effectively during transmission and to ensure that it is sensitive to the appropriate spectral band during the reception interval. (A tuned circuit naturally resonates at a particular frequency, allowing it to absorb energy at that frequency—just as a wine glass does when a singer hits a note that shatters it.) Tuning the coil manually to measure each new NQR spectral line would be a difficult and time-consuming process.

Over the past four years, our research groups have overcome all of these challenges while working to develop a sensitive, low-cost, and reliable NQR-based drug-authentication system. We addressed the issue of low sensitivity with a technique called polarization transfer, which was explored a decade or more ago in the application of NQR to explosives detection. For this, the sample is briefly placed in a magnetic field of about half a tesla, generated by a permanent magnet. The sample is then rapidly removed from the magnet using a motorized actuator. Some complicated physics then ensues between hydrogen nuclei in the specimen (which respond to magnetic fields) and nearby NQR-sensitive nuclei. This process typically increases the NQR signal amplitude by a factor of 5 to 10. At that level, very small samples with only about a gram of active ingredient can be reliably tested within a minute or less. By increasing the magnetic field to 2 T, which can be done readily with rare-earth magnets, we should be able to get even better results, although the stronger magnet would raise manufacturing costs somewhat.

We dealt with RF interference in a few different ways. The first was simply to build proper RF shielding into our equipment. Our prototype also uses adaptive noise cancellation—similar to what's done by noise-canceling headphones—as well as signal-processing methods that can separate the signal from the interference.

To address the awkwardness of having to manually tune the coil for a given frequency band, we used digitally programmable tuning circuits, as have been used in commercial NQR gear in the past. These circuits contain capacitor arrays and miniature relays that are digitally controlled, thus allowing a computer to alter the coil's resonant frequency on command.

We continue to improve our NQR-based drug-authentication system and are also exploring approaches that combine these measurements with near-infrared optical spectroscopy. We hope that such technology might one day see widespread application at hospitals, health clinics, neighborhood pharmacies, maybe even inside some people's homes. Being able to test pharmaceuticals quickly and reliably would go a long way toward improving the integrity and security of the supply chain for essential medicines. ■

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/nqr0919

THE EYES HAVE IT

NEW IRIS-RECOGNITION TECHNIQUES CAN TELL WHETHER AN EYE IS HEALTHY, DISEASED—OR DEAD

By Adam Czajka, Mateusz Trokielewicz & Piotr Maciejewicz

No matter how many times you hold open a cadaver's eyelids to image the irises, each time is uniquely memorable.

One of us (Maciejewicz) once fielded phone calls day and night from mortuary workers at the Medical University of Warsaw's hospital, in Poland. The calls were often placed immediately after a death or the arrival of a cadaver, so that Maciejewicz could get to the mortuary as soon as possible. The bodies would often come in with traces of their final moments: tiny bits of debris from a traffic accident, or electrodes on their skin from failed resuscitations. Maciejewicz may have been focused on imaging the irises, but for him, these traces made every encounter personal, as if he were meeting these people after death. After his work was done, he would thank each person aloud.

The reason why Maciejewicz was scanning eyes in a room full of cadavers was to help answer some lingering questions about the security of iris-recognition systems. As iris scanning starts going mainstream, such questions are becoming more urgent. Around the world, these systems help us skip security lines at many international airports, withdraw money from ATMs, and unlock our smartphones with a glance (Samsung, for example, uses iris scanning, while competitors such as Apple have opted for facial recognition). Governments, including those of Ghana, Tanzania, and Somaliland, have used iris recognition to identify duplicate voter registration records ahead of elections. The world's largest biometric project, operated by the Unique Identification Authority of India, uses iris recognition along with other biometric identification to issue a unique ID number, or Aadhaar, to Indian residents. So far, Aadhaars have been issued to 1.2 billion Indian residents.

Iris recognition is gaining favor for these applications because the iris's structure, like that of a fingerprint, is unique to every individual and doesn't change over the course of one's life. Identical twins, who have genetically identical eyes, do not share the same patterns. Even your own left and right irises are distinct. And while fingerprints are also commonly used to identify an individual, the iris is particularly attractive because it's more complex and therefore more discriminating than other options. In theory, at least, that additional complexity makes it easier to correctly identify an individual and harder to fake someone else's iris.

But it's precisely because of the growing popularity of iris-recognition systems that it's fair to ask how well those systems work. How well can these systems distinguish between a real iris and a replica, such as a high-quality image? Can they recognize an iris that has been—as gruesome as it may sound—plucked from a cadaver? And what about the rare case in which the iris *does* change, because of disease or injury?

Conventional wisdom has held that the iris begins decaying only minutes after death. Thanks to Maciejewicz's work, we discovered that wisdom to be wrong: If cadavers are kept cool, the eyes can still be used for identification for up to three weeks postmortem. Researchers had also assumed that recognition systems could not accurately identify dead eyes, which means that such systems are now vulnerable to exploitation. So to be completely secure, future generations of iris-recognition systems will therefore need more-advanced detection mechanisms, capable of recognizing dead eyes—otherwise, we could find ourselves in a science-fiction future where people could use someone else's dead eyes to access information or locations they're not supposed to. Less fantastical, though, the systems will also need to be flexible enough to adapt in instances where there are changes to an iris due to disease, and precise enough to tell whether the iris in question is a fake.

Today, commercial iris-recognition systems use near-infrared light to illuminate the eye for a scan. Near infrared works well for iris scanning because it is largely not absorbed by melanin, the pigment that among other things determines the iris's color.

Iris-recognition systems also rely on methods to segment parts of the image into iris and non-iris areas. This segmentation allows the system to use an image in which part of the iris is obstructed—for example, by eyelashes, eyelids, or light reflecting off the eye—or when the entire iris is misshapen or damaged, as in more severe cases of postmortem or diseased eyes. Following segmentation, the prevailing approach is to filter the image to make the pattern more pronounced. The pattern is then converted into a binary code. Readily available software can then very quickly hunt for a match between this code and others in a database.

One notable characteristic of these systems is that they all work with a still

PREVIOUS PAGES: ADAM VOORHES/GALLERY STOCK

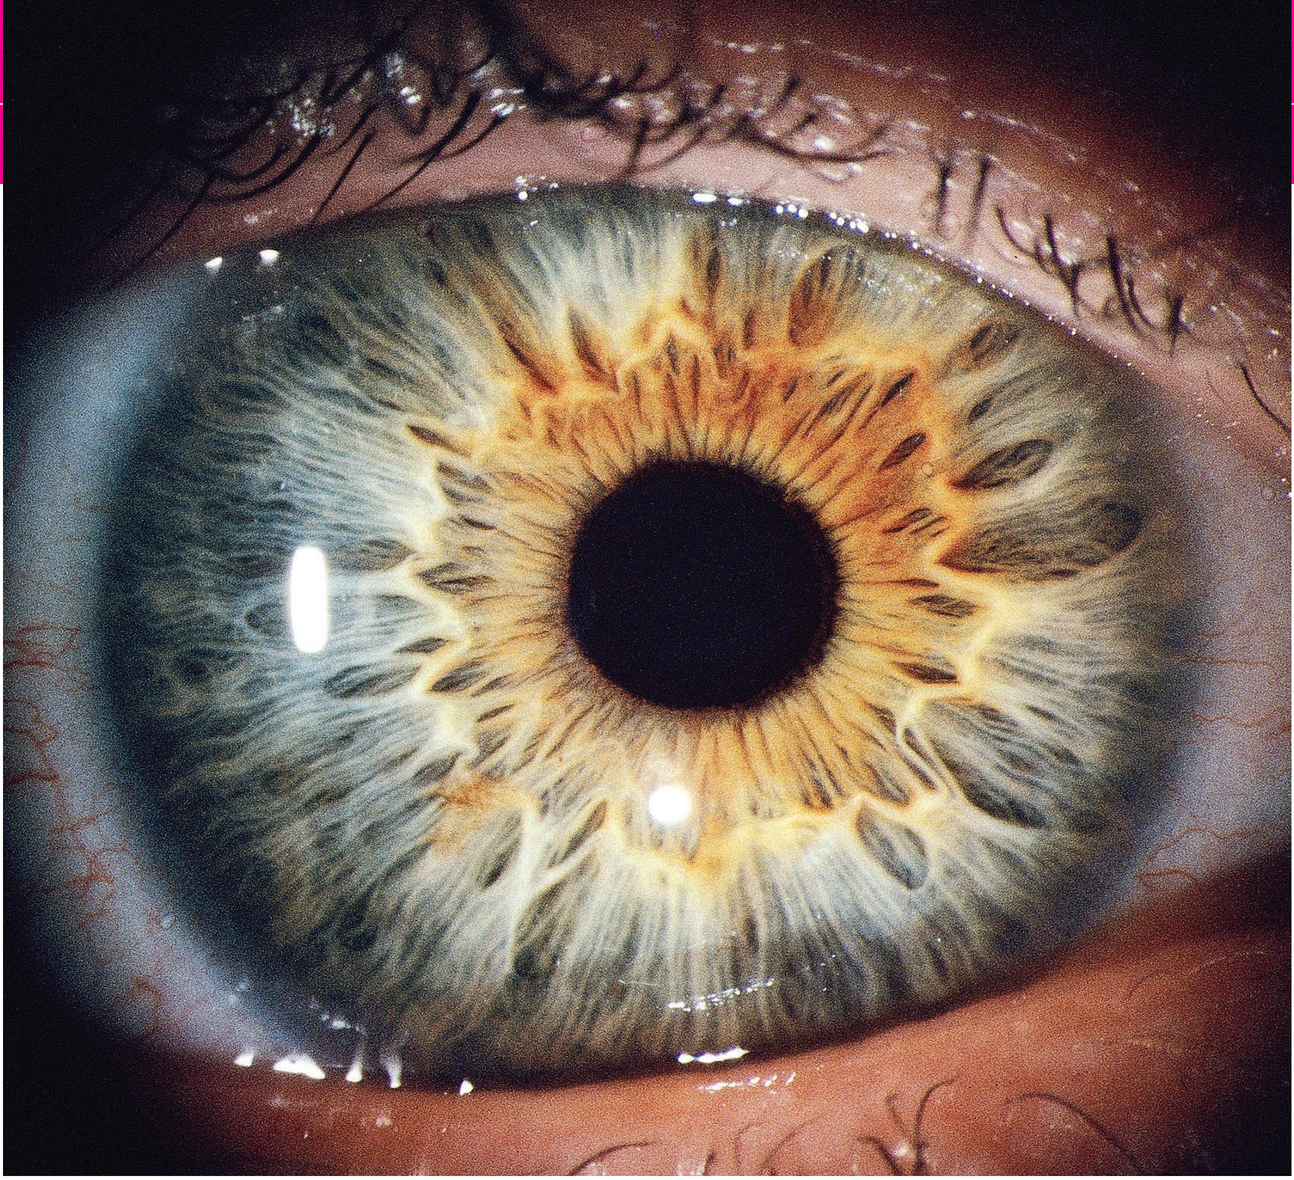

photo of the iris. So a simple test to automatically check whether the eye is "alive"—shining a light on it to see if the pupil contracts—is not always available. And there are other issues. For example, we found in our time at the hospital mortuary that the cornea—that protective, transparent outer layer of the eye—becomes cloudy soon after death. That cloudiness is noticeable in visible light, but near-infrared light basically sees right through it. Furthermore, after death, pupils become fixed in the "cadaveric position," a mid-dilated position similar to the one considered ideal for recognition systems. The cadaveric position makes it tricky to tell a dead eye from a living eye at a glance in normal lighting conditions.

These factors, and our research on iris-recognition algorithms, support our finding that irises remain identifiable up to 21 days after death, a discovery that actually has an important and positive implication: Iris recognition could become a powerful new option for forensic examiners when they need to verify the identity of a corpse. Traditionally, identifications are made from fingerprints, dental records, or DNA, but those options can take hours or even days. Iris scanning, using the same method we used to catalog cadaver irises, could deliver the identification almost immediately.

As we've mentioned, the macabre downside is the possibility of someone using a dead person's eyeball to gain access to secure locations or information. Sci-fi movies have already trotted out such grisly scenarios, but it's never happened in the real world. But if iris-recognition systems are more vulnerable than expected, it's reasonable to ask how else they might be tricked. There have been plenty of instances in which more mundane schemes have been used to get past these systems.

EYE OF THE BEHOLDER: The iris of the human eye has tremendous detail, and every person's irises are unique, even in twins. Structures such as the crypts of Fuch (the oval-shaped openings surrounding the iris's outer rim) are complex enough to serve as clear markers of identity.

JOE MCNALLY/GETTY IMAGES

Back in 2002, a group of German researchers and hackers demonstrated that the sensors in commercial recognition systems could be tricked by a person holding up to the scanner a paper printout of a photo-

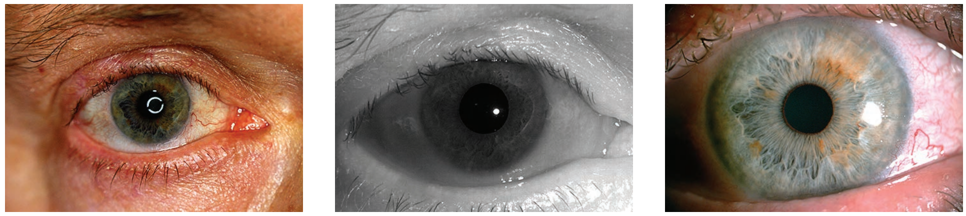

CLEAR EYES AHEAD: These three images of the same eye were taken with a Canon EOS 1000D camera [left]; an IrisGuard camera, which uses infrared to better capture the iris's fine detail [center]; and a Topcon ophthalmic device [right], which is used during eye exams.

graph of an iris, with a hole in place of the pupil. Despite that revelation, recognition systems remain vulnerable to this very trick. In 2017, Chaos Computer Club, which bills itself as Europe's largest association of hackers dedicated to making the public aware of data security issues, showed that the iris scanner on the Samsung Galaxy S8 could be deceived by using a photograph of an iris with a contact lens laid on top of it. The one catch is that the camera used to take the photo must be capable of capturing near-infrared light. Digital cameras typically have a filter that removes this light before it reaches the image sensor, but for many digital single-lens reflex (DSLR) cameras it's not difficult to remove the filter.

Successful biometric attacks like the Galaxy S8 demonstration have inspired security researchers to step up their work on detecting and handling such attacks. Those efforts have resulted in a host of effective countermeasures.

One possible measure relies on the use of photometric stereo, a computer-vision technique that captures the three-dimensional structure of an object. The technique works by illuminating the object from multiple directions, one direction at a time, and photographing the results. It's possible, using the resulting collection of images, to determine the orientation of the object's surface at any point by comparing how light reflects from different angles. We showed earlier this year that photometric stereo can detect when someone is wearing contact lenses with someone else's iris pattern, which might bamboozle an ordinary iris-recognition system. Fortunately, most modern iris-recognition systems can be adapted to use photometric stereo without swapping out hardware. Another possibility is that systems could be adjusted to spot telltale anomalies that printers leave on printed images.

Several techniques could help these identification systems detect a dead eye. One is to add an additional, thermal sensor to the system in order to detect an eye too cold to be part of a living person. Adding an additional sensor would be a relatively expensive proposition, however.

It's also possible to tell the system to pay attention to the pupil as well as the iris. In a healthy, living eye, the pupil contracts and dilates in response to changes in lighting. With that involuntary reaction in mind, we built a general model of how a human pupil reacts to changes in brightness. Our goal was to develop a method to verify whether an iris is alive or not based on that reaction.

The method we developed images the eye multiple times to see whether it's actually responding to changes in brightness; the process takes about 3 seconds to verify the iris is alive. For comparison, even iris-recognition systems that require just one snapshot usually take about that much time to make an identification. And there's no reason why the two techniques couldn't be used in parallel. That means it would be possible to use our method to verify that an iris is part of a living eye while also confirming identity with a still image, with no additional time required.

Let's say, however, that you want a recognition system that can quickly detect whether or not an iris is dead using only a single snapshot. For that, you'll have to turn to convolutional neural networks, a subset of machine learning geared mainly toward analyzing images. In 2018, we developed a convolutional neural network with reference images of living and dead irises, and let it determine for itself what set the two categories apart. After approximately 2 hours of training, the network could correctly decide whether an image showed a live iris or a dead one 99 percent of the time.

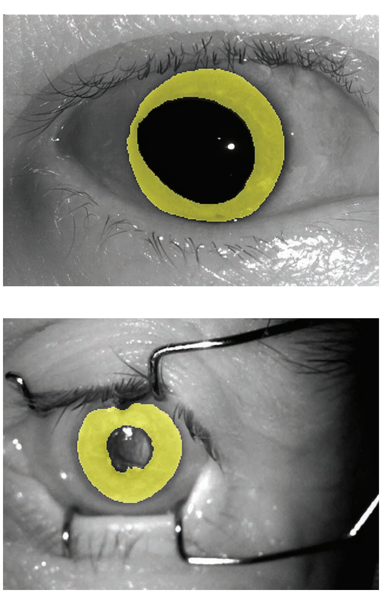

DEAD-EYE ACCURACY: To identify an individual, iris-recognition systems use segmentation to separate the iris and non-iris portions of the image. Machine learning can accurately segment both diseased eyes [top] with distorted irises and dead eyes [bottom], which may have begun to decay.

ADAM CZAJKA, MATEUSZ TROKIELEWICZ & PIOTR MACIEJEWICZ (8)

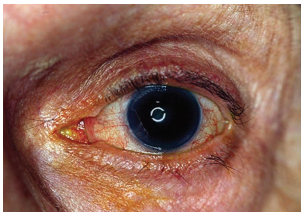
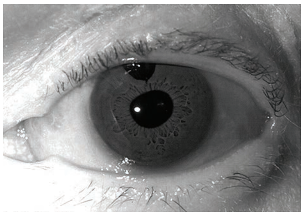
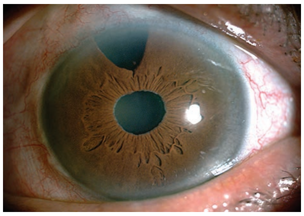

CLOUDY EYES AHEAD: In contrast with the iris on the opposite page, this one isn't healthy. That might not be as obvious when imaged with the Canon [left], but the IrisGuard and Topcon images [center and right] effectively capture the deformed iris structure.

So there's reason for optimism that iris recognition will rise to future security challenges. The ability to recognize iris patterns even after death may break new ground in forensics. In addition, the use of video to record the iris's reflexive response to light could form the basis of a robust recognition system immune to spoofing schemes that make use of cadaver parts or detailed fakes.

But one big problem remains: A person's iris pattern can undergo changes because of disease. These changes can be significant enough to render an iris-scanning system unable to recognize the iris.

Several different ocular diseases can alter the pattern of the iris. Rubeosis iridis, iridodialysis, and synechiae all distort the shape of the iris and the pupil. Pterygium, bacterial keratitis, and hyphema make an iris's pattern less visible. All of these conditions, we found, can wreak havoc with iris-recognition systems. They can change the iris so much it can't be recognized, of course. But they can also make it difficult or impossible to record a person's pattern because the shapes have been distorted beyond the system's ability to accommodate them.

Iris-biometrics researchers are still debating how best to address recognition failures due to medical conditions. The U.S. National Institute of Standards and Technology (NIST), for example, hosted a meeting this past June for its Iris Experts Group to discuss the problem. One partial solution is to have people re-enroll their eyes in recognition systems after any medical treatment that might alter the iris. However, this solution doesn't get at the underlying problem, which is that medical conditions make iris patterns harder to record because such conditions blur and distort the iris's features.

We began studying the effects of ocular diseases on iris-recognition systems in 2013. We found that in many cases where disease has made recognition impossible, it's because the recognition system had erroneously segmented the image of the iris. (As a reminder, segmentation is the technique that recognition systems use to separate the iris and non-iris portions of the eye so that the pattern of the iris can be properly identified.) The problem isn't that the system can't spot the pattern; it can't distinguish between what is and what isn't an iris in the first place.

Essentially, all iris-recognition software written to date assumes that the shape of the iris is circular, or nearly so. That assumption is typically true, and it speeds up and simplifies image processing. However, some eye diseases can change the shape of the pupil or of the outer rim of the iris, or both, leaving the overall shape unrecognizable to a scanner.

The best way to make iris-recognition systems more accessible to more people is to relax the assumptions these systems make about the iris's shape. Remember how we used convolutional neural networks to distinguish between live eyes and dead ones? We and other researchers are using the same techniques now to create systems whose assumptions are more flexible about what the iris should look like. What we're all striving for is an all-in-one system that can distinguish among healthy, diseased, and dead eyes reliably.

Another intriguing future possibility are recognition systems that would identify an individual based on minute details of the iris, rather than the iris pattern as a whole. This technique would be similar to how forensic scientists match fingerprints today, by comparing just a few points of interest rather than the entire whorl. Crypts, for example, are minuscule holes in the iris tissue, usually located near the pupil, that adjust as the pupil changes shape. It may be that an individual's crypts are as unique as the ridges of their fingerprints.

And if crypts don't work out, there are other options. In a study carried out in 2017, we asked participants to examine two iris images and tell us if they were the same iris or not. While our participants scrutinized the images, we tracked their own eye movements to see where they looked. We've scoured these gaze maps, as they're called, to identify the spots people tended to focus on. The next step for us is to direct convolutional neural networks to focus on the same locations as they attempt to match an iris scan.

Iris-recognition systems have been around for 25 years, but only now are we and others finally addressing some longstanding flaws. In so doing, we are advancing our knowledge about the iris. By better understanding the assumptions we've made for these systems, and improving our techniques for working around those assumptions, future recognition systems will capture more information in the blink of an eye. ■

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/irisbiometrics0919

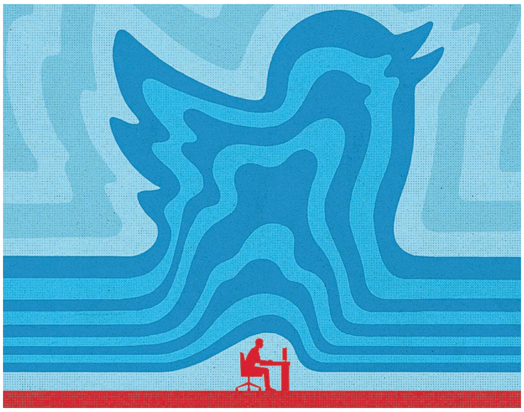

DAN PAGE/THE I SPOT

CONTINUED FROM PAGE 18 |

Making real-world connections: Following and engaging with people you find fascinating on social media can sometimes lead to real-world connections.

When Megan Killian, a biomechanist at the University of Delaware, was looking to use some tools from the field of neurophysiology for her experiments, she started following and interacting with several neuroscientists on Twitter. One of them offered to help her out with her research. Ultimately, they applied for and won a grant together. "If it weren't on Twitter, I don't think I would have crossed paths with him," says Killian.

Twitter can also make it easier to make friends at conferences. Whenever Killian goes to a conference, she puts out feelers on Twitter, asking who is going to be there. It opens up opportunities for her to make some virtual connections before she heads to the meeting. Killian also uses another trick to ignite conversations. She walks up to people whom she recognizes from the site and breaks the ice by saying, "I follow you on Twitter."

Humanizing yourself: Even people who followed you to get updates on your work appreciate a peek into your personality once in a while. "I do think there's benefit to being more human on Twitter than just a bot that puts out 'Our paper came out recently' and 'We recently got this grant,' " says Killian. "Those announcements are great and should definitely be celebrated. But I think there's a really great opportunity to use social media like Instagram, Twitter, and even Snapchat to connect with people on a personal level." For example, don't hesitate to share some slice-of-life posts about academia.

However, Dietterich says it's wise to stay focused. "If you want to engage in political commentary or post vacation photos, create separate accounts for those (or maybe go to different platforms)," he says.

—DINSA SACHAN

↗ **POST YOUR COMMENTS** at https://spectrum.ieee.org/social0919

CRUEL BEAUTY

For more than a century, fashion-conscious consumers have relied on the electric hot comb to style their hair—to straighten and smooth coarse hair, or add a wave or curl to fine hair. Much like the clothes iron, the hot comb predated electrification. The original combs, made from ceramic or metal, were heated in a water bath or over a gas burner. The 20th-century electric version provided heat on demand and slightly better temperature regulation. Unfortunately, electricity could not prevent the occasional burned scalp or the pungent smell of singed hair. Sometimes the pursuit of beauty can be painful. ■

↗ **For more on the history of electric hot combs, see https://spectrum.ieee.org/pastforward0919**

SMITHSONIAN NATIONAL MUSEUM OF AFRICAN AMERICAN HISTORY AND CULTURE

南方科技大学
SOUTHERN UNIVERSITY OF SCIENCE AND TECHNOLOGY

Professor/ Associate Professor/ Assistant Professor
School of Microelectronics
Southern University of Science and Technology (Shenzhen, China)

Post Specifications

School of Microelectronics (SME) National Exemplary School of Microelectronics, Southern University of Science and Technology (SUSTech) invites highly qualified candidates to fill multiple tenure-track/tenured faculty positions in the areas of (but not limited) Emerging Microelectronic Devices (Wide-bandgap, Nonvolatile memory, MEMS Sensor), and IC-Chip Designs (Future Computing/Communication/ Biomedical SoC).

Junior applicants should have (i) a PhD degree in related fields; and (ii) outstanding potential in teaching and research. Candidates for senior post are expected to have demonstrated exceptional academic leadership and strong commitment to be excellent in teaching, research, and services.

Applications

Submit (in English, PDF version) a cover letter, a statement in research and teaching, a CV plus copies of 3 most significant publications, and contacts of three referees to: sme-hr@sme.sustech.edu.cn entitled with “Apply for Faculty Position”. Applicants are required to specify the rank of the position in their letter of application. The positions will be open until they are filled by appropriate candidates.

For more information, please visit h**ttp://ohr.sustc.edu.cn/sustczp/product/recruit/a.do?action=toZPGWList2&entityId=T_RECRUIT_PLAN&postType=2&selectedId=1009881**.

Salary and Fringe Benefits

Salary and startup funds are highly competitive, commensurate with experience and academic accomplishment. All regular faculty members will join the tenure-track system in accordance with international practice for progress evaluation and promotion. Applicants are encouraged to check out the details about the university at: **http://www.sustech.edu.cn**.

Micron Technology, Inc., is seeking the below positions for its semiconductor R&D facility in Boise, ID; its manufacturing facility in Manassas, VA; and its sales and design facilities in Folsom and Milpitas, CA; and design facilities in Austin, TX and Longmont, CO; and worksites in Irvine, CA; Novi, MI; and Allen, TX.

The following Micron subsidiaries are also seeking positions: Micron Semiconductor Products, Inc., at its headquarters in Boise, ID; and sales facilities in Meridian, ID; and Folsom and Milpitas, CA.

Electrical, Electronics, Communications, Chemical, Industrial, Mechanical, Materials, Computer System Analysts, and Software Engineering; Physics, Materials Science, Engineering Manager and other related Engineering occupations. Marketing, Sales, Logisticians, Finance, Accounting, and other related business positions.

Please submit your resume online:
http://www.micron.com/jobs

Resume and/or cover letter must reflect each requirement or it will be rejected. Upon hire, all applicants will be subject to drug testing/screening and background checks.

Note: Some of these positions may require domestic and international travel for brief business purposes. Please read the full job description when applying online for such requirements.

EOE

Big Data Warehouse Integration Specialist - (Regions – Hoover, AL) Multiple positions- Utilize minimum requirements to analyze, design, develop, test complex data transformation solutions for Enterprise Data Warehouse system and architect scalable and performing ETL solutions. Develop ETL programs using Informatica and data stored in DB2, SQL Server, VSAM, ISAM, Oracle, flat files. Write / deploy stored procedures, tune SQL scripts/queries and use database monitoring tools. Build ETL solutions to extract and transform data from mainframe systems and processes meeting data lifecycle management standards. Design systems that are recoverable, auditable, parameter driven. Work with Structured and Unstructured data. Implement ETL solutions with Processing and Error Handling techniques. Minimum Requirements: Must have a Bachelor’s degree or foreign degree equivalent in Computer Science, Information Systems, or related field and 5 years of experience in the job offered or 5 years of experience in a data warehouse development position. Experience may have been gained concurrently and must include 5 years of experience in each of the following: Informatica PowerExchange and PowerCenter tools and architecture OR Syncsort DMX-h tools and architecture; Oracle databases procedures, structures and data manipulation; Unix or LINUX scripting; data warehouse modeling concepts; ETL standards, methodologies, guidelines and techniques; DBMS concepts of data structures and normalization. Must have legal authority to work in U.S. Resume to Regions Bank, Attn: Ethan Dorman, Talent Acquisition Partner, Riverchase Center, 2090 Parkway Office Circle, Hoover AL 35244.

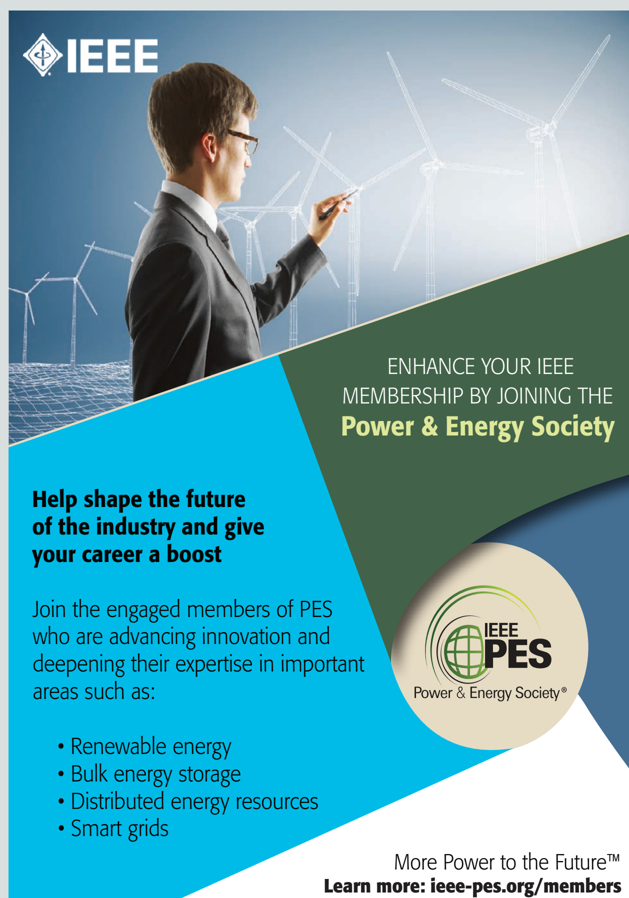


70
70
70
70
70
70