ASIAN SCIENTIST
Issue 05
January 2019

SUPERCOMPUTING ASIA

SUPPORTING SCIENCE AT ALL SCALES

THE SILICON BEHIND THE SCIENCE

+

5 WAYS SUPERCOMPUTERS HELP ADVANCE SCIENCE

QUANTUM COMPUTERS ON THE HORIZON

ALL EYES ON THE FUTURE

MCI (P) 127/04/2018

CONTENTS

Issue 05
January 2019

FEATURES

p. 22
Quantum Computers on the Horizon
Not if, but when

p. 26
All Eyes on the Future
Bringing sight to machines

p. 30
Supercomputing Shouldn't Be Rocket Science
Putting computing power in the hands of scientists

p. 32
Embracing AI in Singapore
How the little red dot is getting AI-ready

p. 34
Inspiring Minds at SC18
Highlights from Dallas, Texas

COVER STORY
p. 16
SUPPORTING SCIENCE AT ALL SCALES
The silicon behind the science

p. 6
Digital Dispatch
Supercomputing news from around the world

p. 10
5 Ways Supercomputers Help Advance Science
The data-crunchers behind complex analyses

p. 36
Business Bytes
The latest industry moves

p. 38
Super Snapshot
Moore's law: Beyond the wall

EDITOR'S NOTE

While the scientific method remains fundamentally unchanged, technological advances over the past few decades have transformed how scientific research is carried out. Today, just about every field of science is more data-intensive, collaborative and swift-moving than ever before.

In this issue, we explore how scientists are working on phenomena at every imaginable scale. From the astronomical distances of cosmology to the subatomic realms of quantum mechanics, scientists are using supercomputers to get the most out of their data (*Supporting Science At All Scales*, p. 16). We also consider five of the thorniest problems in science that supercomputers are now poised to help scientists tackle (*5 Ways Supercomputers Help Advance Science*, p. 10).

But all the supercomputing resources in the world will not drive scientific advances unless they are made readily accessible to scientists. In *Supercomputing Shouldn't Be Rocket Science* (p. 30), we take a look at tools and initiatives that aim to place the power of supercomputing at scientists' fingertips.

Finally, we turn our attention to quantum computing, a field that not only holds promise for taking computing beyond the physical limits of Moore's law, but also for helping scientists describe natural phenomena more thoroughly and accurately. Find out more in our interview with IBM's Dr. Christine Ouyang (*Quantum Computers on the Horizon*, p. 22).

Rebecca Tan, Ph.D.
Editor-in-chief
Supercomputing Asia

/asianscientist @asianscientist
Asian Scientist Magazine

www.asianscientist.com

SUPERCOMPUTING ASIA

EDITORIAL ADVISORY COMMITTEE
Prof. Tan Tin Wee
Prof. Satoshi Matsuoka
Prof. John Gustafson
Yves Poppe

CEO & PUBLISHER
Dr. Juliana Chan

EDITOR-IN-CHIEF
Dr. Rebecca Tan

SENIOR EDITOR
Dr. Sim Shuzhen

DIGITAL EDITOR
Dr. Jeremy Chan

CONTRIBUTORS
Jerry Lim
Cindy Lim
Brenda Lau

SENIOR DESIGNERS
Lam Oi Keat
Muhamad Azlin

SALES & MARKETING
Samantha Yeap

PUBLISHED BY
Wildtype Media Group Pte Ltd

DISTRIBUTED BY
Pansing Distribution Pte Ltd

HEAD OFFICE
Wildtype Media Group Pte Ltd
71 Ayer Rajah Crescent #07-03
Singapore 139951
Tel: 65-6909-2279
Website: www.asianscientist.com/subscribe

WILDTYPE MEDIA

Combining savvy communication with technical rigor, Wildtype Media Group is Asia's leading STEM-focused media company, spanning digital, print, custom publishing and events.

Brands under Wildtype Media Group include the flagship *Asian Scientist Magazine* and *Supercomputing Asia*, award-winning titles available in print and online. Through its Asian Scientist Intelligence consulting division, Wildtype Media Group provides a full suite of marketing and communication services to industry and academic clients hoping to engage Asia.

www.wildtype.media

To order reprints or e-prints, or request permission to republish *Supercomputing Asia* content, please contact hello@wildtype.media

Digital Dispatch

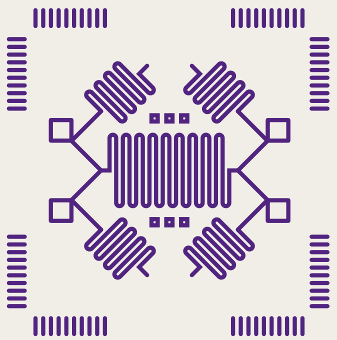

CHINA INVESTS US$145 MILLION IN SUPERCONDUCTING-BASED COMPUTING

The Chinese Academy of Sciences (CAS) is spending one billion yuan (US$145 million) on efforts to build a superconducting computer, with a working prototype targeted for 2022, reported the *South China Morning Post*.

With near-zero electrical resistance, superconducting circuits promise orders of magnitude higher energy efficiency and processing speed than their semiconducting counterparts, but high costs and technical complications have so far precluded their use in supercomputers.

According to the *South China Morning Post* report, Chinese scientists have already developed new integrated circuits using an unnamed superconducting material, tested an industrial process for the mass production of superconducting chips, and nearly completed the design of the superconducting computer's architecture.

"Superconducting digital circuits and superconducting computers... will help China cut corners and overtake [other countries] in integrated circuit technology," noted CAS president Professor Bai Chunli during a visit to the Shanghai Institute of Microsystem and Information Technology.

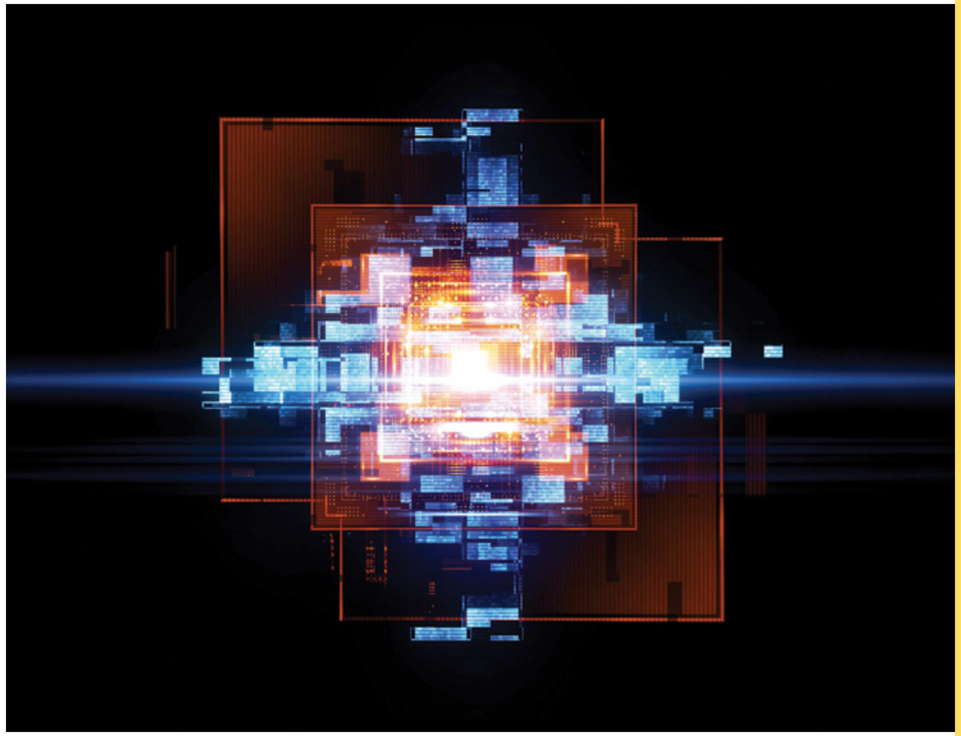

CHINA LAUNCHES THIRD EXASCALE PROTOTYPE

China has unveiled its third exascale supercomputer prototype, further improving the country's odds of meeting its 2020 exascale target, reported Xinhua News Agency. The new prototype, developed by Dawning Information Industry (also known as Sugon), is for a machine called the Shuguang, which is expected to be installed at National Supercomputing Centers in Shanghai and Shenzhen.

The Shuguang model follows two other exascale prototypes already revealed by China in 2018, the first for the Tianhe-3 machine to be housed at the National Supercomputing Center in Tianjin, and the second for the Sunway system at the National Supercomputing Center in Jinan. While no official details are yet available, the Tianhe-3 is expected to be an ARM-based supercomputer utilizing Chinese semiconductor company Phytium's Xiaomi platform, and the Sunway system is expected to run on an updated version of the ShenWei processor currently powering the Sunway TaihuLight.

Meanwhile, *Top500 News* predicts that the Shuguang will be based on x86 processors, as Chinese chipmaker Hygon now has access to the technology via a licensing agreement with US semiconductor company AMD.

US CONGRESS GIVES QUANTUM COMPUTING A BILLION-DOLLAR BOOST

The US House of Representatives has passed the National Quantum Initiative Act (NQIA), a national strategy to accelerate research and development in quantum computing. The ten-year initiative is split into two five-year plans, the first of which will receive US$1.275 billion in funding over fiscal years 2019 to 2023.

Involving the National Institute of Standards and Technology, the National Science Foundation (NSF) and the Department of Energy (DOE), the NQIA will invest in fundamental quantum information science and technology research, demonstrate applications in quantum technology, and support a quantum information science and technology workforce. The DOE, for example, will receive US$625 million to set up five National Quantum Information Science Research Centers, which will facilitate collaborations across national laboratories, universities, research institutions and the private sector.

The nod from Congress is seen as a bid to maintain US leadership in quantum computing and other technologies, as countries like China, Japan, Russia and the UK ramp up funding for quantum research.

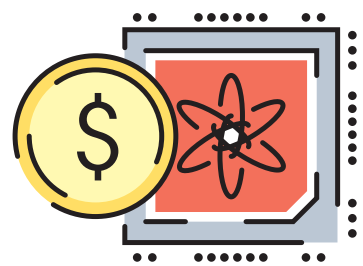

EU-DEVELOPED CHIPS TO PAVE THE WAY TO EXASCALE

The European Union has launched the European Processor Initiative (EPI), an ambitious strategy to develop low-power microprocessors for the domestic market and future exascale supercomputers.

"The EPI is an important step of a strategic plan to develop an independent and innovative European supercomputing and data ecosystem, and will ensure that the key competence of high-end chip design remains in Europe, a critical point for many application areas," said the European Commission's Mr. Andrus Ansip, vice president and commissioner for Digital Single Market, and Ms. Mariya Gabriel, commissioner for Digital Economy and Society, in a statement.

First-generation chips are expected to power the EU's pre-exascale supercomputers by 2020, and will be followed by second-generation chips for exascale systems from 2023 onwards. The chips are likely to be based on the ARM and RISC-V processor architectures for general-purpose processing and throughput acceleration respectively, reported *Top500 News*.

Digital Dispatch

CHINA'S SENSETIME SIGNS MOU WITH SINGAPORE PARTNERS

Chinese artificial intelligence (AI) unicorn SenseTime has signed memoranda of understanding (MOUs) with three Singapore organizations: the National Supercomputing Centre (NSCC) Singapore, Nanyang Technological University (NTU) and Singtel. The parties will leverage each other's strengths and customer bases to advance AI research and accelerate digitization in Singapore and Asia.

"As Singapore serves as SenseTime's international hub, we're committed to building an AI ecosystem with our local partners and servicing our customers with leading AI technologies. We look forward to promoting the development of AI in the ASEAN region with their partnerships," said Mr. Martin Huang, managing director of SenseTime Singapore.

"With this partnership, we will build upon and extend NSCC's existing petascale distributed HPC infrastructure. We hope to achieve and bridge the HPC ecosystem by setting up a GPU-accelerated, AI-centric, supercomputing infrastructure and platform," said NSCC Singapore chief executive Professor Tan Tin Wee.

TSINGHUA TEAM WINS INAUGURAL ASIA PACIFIC HPC-AI COMPETITION

A student team from Tsinghua University, China, took first place in the inaugural Asia Pacific High Performance Computing-Artificial Intelligence (HPC-AI) Competition held in Singapore. The event, co-organized by the HPC-AI Advisory Council and the National Supercomputing Centre (NSCC) Singapore, saw 18 university teams from seven countries compete to develop new solutions for weather forecasting and image recognition.

Teams were tasked with achieving high simulation speeds using popular weather forecast tools, with the aim of reducing forecasting time while maintaining precision. They were also required to optimize an AI framework for image recognition, which is relevant for smart city- and security-related applications. The second and third prizes were awarded to students from Taiwan's National Cheng Kung University.

"If you think that supercomputing competitions are but an exercise conducted in an ivory tower, offering no practical benefits to mankind—think again... the winning solutions show clearly how supercomputing can align research, university work and commercial considerations to drive innovation that benefits all," said Professor Tan Tin Wee, chief executive of NSCC Singapore.

WHAT'S UP!

COMING YOUR WAY: SUPERCOMPUTINGASIA 2019

Asia's premier HPC conference SupercomputingAsia (SCA) will be back from March 11–14, 2019 in Singapore. In 2018, the event was attended by over 800 delegates from more than 24 countries.

Organized by the National Supercomputing Centre Singapore, SCA19 will focus on 'HPC Futures—Hyperscalers, Exa, AI, Quantum and Beyond,' and will again feature insightful keynotes from global HPC leaders in academia and industry, including Professor Lu Yutong, director of the National Supercomputing Center, Guangzhou, China, and Dr. Christine Ouyang, distinguished engineer and master inventor at IBM.

Winners of the inaugural Move That Data! Data Mover Challenge will be announced at SCA19, and attendees will be able to witness a live demonstration of the winning software deployed across servers in Australia, Japan, Singapore and the US.

Other SCA19 program highlights include the ASEAN HPC Workshop, the Conference on Next Generation Arithmetic (CoNGA), as well as talks on green data center strategies, precision medicine, quantum computing and more. Save the date!

For more information, visit https://www.sc-asia.org/

WHAT

SupercomputingAsia 2019

WHEN

March 11–14, 2019

WHERE

Singapore

DON'T MISS ISC HIGH PERFORMANCE 2019

This year's edition of ISC High Performance—the world's oldest and Europe's most important conference and networking event for the HPC community—will be held from June 16–20, 2019 in Frankfurt, Germany.

First held in 1986, ISC High Performance brings together engineers, IT specialists, vendors, scientists, students and other members of the HPC global community. The 2019 conference is expected to draw over 3,500 researchers and commercial users, along with 150 exhibitors.

Participants can look forward to thought-provoking discussions centered around 13 topics selected by the ISC 2019 program team, including next-generation high performance components, exascale systems, blockchain and cryptocurrencies, parallel processing in the life sciences, and machine learning and big data analytics. Registration will open in spring 2019.

For more information, visit https://www.isc-hpc.com/

WHAT

ISC High Performance 2019

WHEN

June 16–20, 2019

WHERE

Frankfurt, Germany

5 WAYS SUPERCOMPUTERS HELP ADVANCE SCIENCE

Supercomputers are helping scientists simulate scientific processes, make predictions and create data-based solutions, tackling the biggest scientific questions one calculation at a time.

By **Brenda Lau**

The data-crunchers behind complex analyses

Supercomputers are the heavyweight champions of the computer world—China's most powerful supercomputer, the Sunway TaihuLight, for example, can perform 93,000 trillion calculations per second, about a million times faster than a typical laptop computer.

This astronomical processing power makes supercomputers perfect for tackling some of the biggest scientific questions researchers grapple with today. With the machines' capability to analyze large amounts of data with multiple variables, they may put researchers one step closer to long-sought goals—understanding the origins of the universe, perhaps, or mapping out the human brain.

Illustrations by Lam Oi Keat/*Supercomputing Asia*

MAPPING OUT THE UNIVERSE

Tracking the births and deaths of stars and galaxies can prove challenging using just observational astronomy—even the best telescopes have limitations.

Here's where supercomputers can help: by feeding supercomputers with data from satellite observations, theoretical astronomers can obtain a virtual simulation of the universe. Currently, it is possible to simulate our universe—from star formation to the distribution of dark matter—in a cube-shaped volume with sides nearly one billion light-years long. However, with the ATERUI II, a Cray XC50 supercomputer at the National Astronomical Observatory of Japan (NAOJ), the dimensions of such simulations can be extended up to 5.4 billion light-years on each axis.

With 40,200 cores linked to achieve a theoretical peak performance of 3.087 petaFLOPS, the ATERUI II is the fastest available parallel supercomputer for astrophysical simulations, and can tackle problems too difficult for previous computers. For example, it can calculate the mutual gravitational forces between the 200 billion stars in the Milky Way, instead of separating the stars into smaller groups like other simulations do. As such, ATERUI II will generate a full-scale, high-resolution model of the Milky Way.

"The age of the universe is 13.8 billion years," said Professor Eiichiro Kokubo, project director for the Center for Computational Astrophysics at NAOJ. "With ATERUI II, we could explore the universe, from the past to the future."

FINDING DRUGS TO FIGHT BUGS

There is an urgent need to find new and better treatments for infectious diseases such as Zika, dengue and tuberculosis. The problem is that researchers are often slowed down by the sheer number of candidate compounds and variables to consider.

IBM has pioneered a crowdsourced 'supercomputer' that uses the idle time of computers and smartphones around the world to address this problem. Running at a collective speed of 1.1 petaFLOPS, the World Community Grid (WCG) enables researchers to screen small-molecule drugs and millions of drug-like compounds from existing databases against models of microorganisms and their proteins. These results are shared in the public domain and made available to the scientific community as part of WCG's commitment to improve human health and welfare.

The compounds that show the most promise are then tested in the laboratory. Instead of having to screen compounds over years or even decades in the laboratory, WCG helps perform these initial tests in a matter of months.

Over two dozen projects have been powered by WCG in the past 13 years, and as of March 2018, the platform boasts a combined total run time of 1.5 million years. During that time, 20 projects have been completed, yielding important discoveries such as a protease inhibitor that disables the replication of the dengue and West Nile viruses, as well as several molecules that are effective against malaria and drug-resistant tuberculosis.

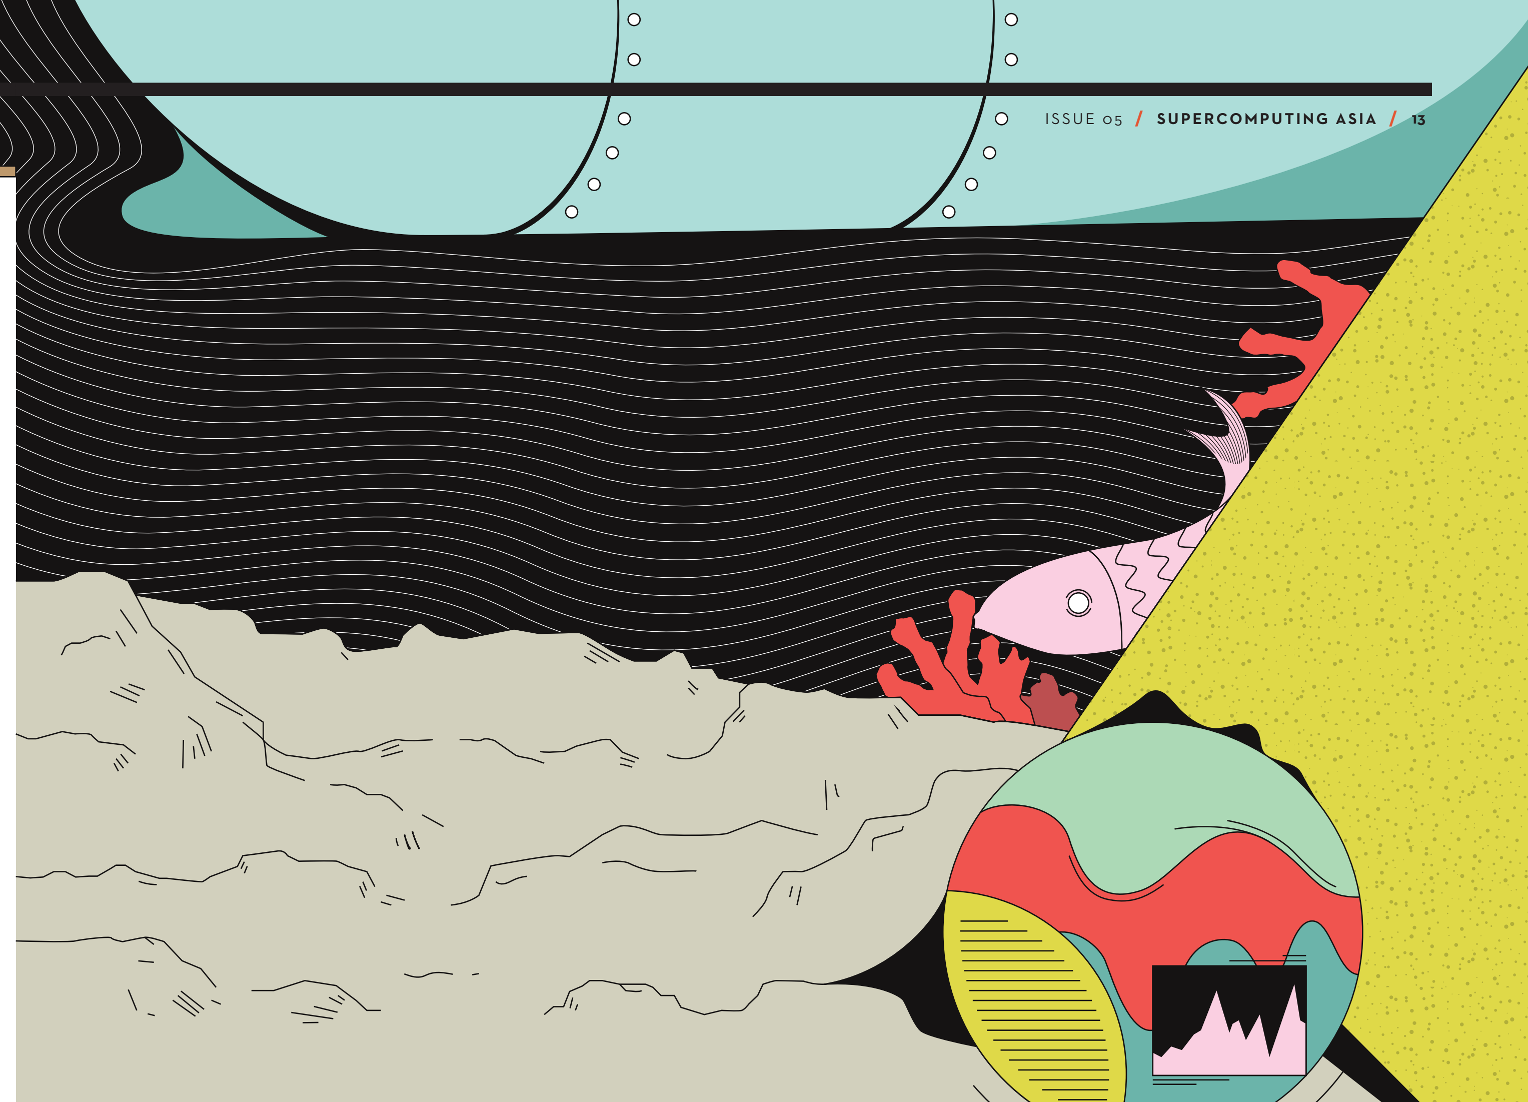

SIMULATING THE SEAS

Climate change affects the oceans as much as it affects the atmosphere. The marine environment is already registering the effects of the 0.7 degree Celsius increase in global temperatures the world has experienced since pre-industrial times. Key changes include the reduction of plankton (microorganisms that form the basis of marine food chains), coral bleaching and damage to marine ecosystems, among others.

Supercomputers could help scientists better understand the impact of these changes on the oceans. China, for example, is building an exascale supercomputer—potentially the world's first—to support ocean research. The new machine is expected to run ten times faster than the 93-petaFLOPS Sunway TaihuLight, and will most likely be based on the coast of Shandong province.

With an expected completion date as early as 2020, the supercomputer will process data on the world's oceans, collected from Chinese vessels, naval outposts and unmanned monitoring facilities—including a global network of buoys, sea floor sensors, underwater gliders and satellites. These data contain a wealth of information, including levels of trace chemicals, sea current readings, regional weather data and water density anomalies.

The comprehensive analysis of such vast amounts of data will allow researchers to simulate the oceans with unprecedented resolution, and hence make more reliable forecasts of important climate-related phenomena such as El Niño. Ultimately, this information will help scientists and policy makers make better-informed decisions on vital issues such as cuts in greenhouse gas emissions.

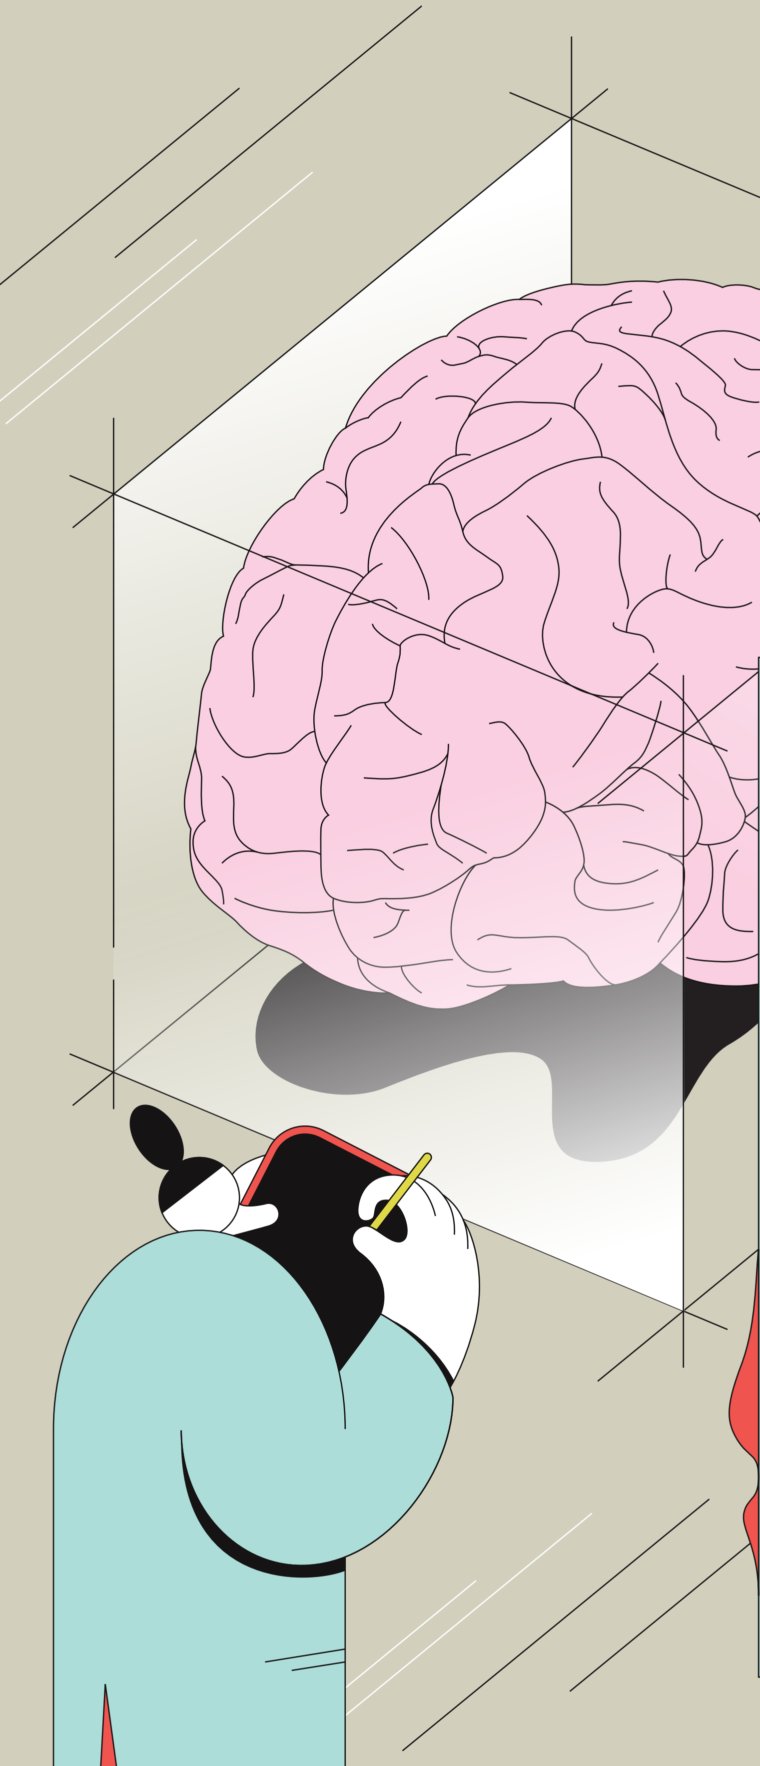

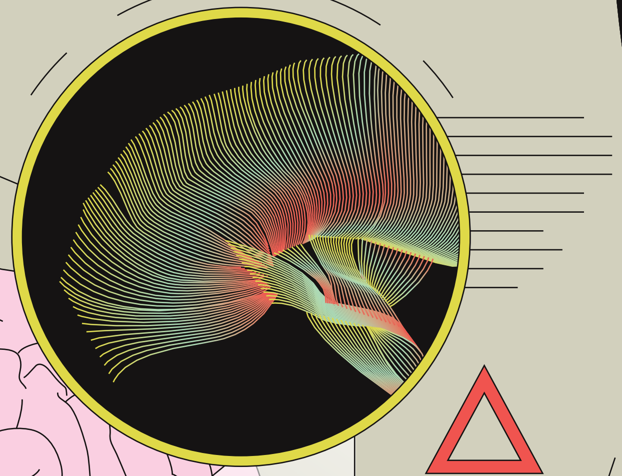

MAPPING THE HUMAN BRAIN

Current biology textbooks will tell you that the brain works based on the firing of neurons and the movement of neurotransmitters across synapses. But how is all this hardware wired together?

Constructing a map of the brain could take the lifetimes of many scientists, as the brain has more connections than the Milky Way has stars. But thanks to supercomputers, this timeframe will soon be dramatically shortened. An exascale supercomputer named Aurora 21 is expected to be deployed at the Argonne National Laboratory in the US by 2021, running one quintillion operations in parallel—roughly similar to the processing power of the human brain.

Scientists plan to use the Aurora 21 to piece together millions of two-dimensional images and reconstruct the entire brain in three dimensions. This will create a comprehensive map of neural connections, known as a connectome.

Instead of stopping at one brain, scientists could also map multiple brains and compare them to learn about how connectomes vary between different individuals—between an adult and a baby, for example, or between adults trained in different skills. Such analyses could also could shed light on the as-yet mysterious processes that underpin human learning, behavior and psychiatric disorders.

RECREATING EARTHQUAKES

Earthquake prediction and simulation is an inexact science. But scientists hope to use supercomputers to address the myriad variables at play during geologic events, leading to better prediction and preparedness.

For example, by developing software to efficiently process 18.9 petaFLOPS of data on the Sunway TaihuLight supercomputer, Chinese researchers were able to recreate three-dimensional visualizations of a devastating earthquake that occurred in Tangshan, China in 1976.

The simulations incorporated data covering the entire spatial area of the quake—a sprawling zone measuring 320 kilometers by 312 kilometers and running 40 kilometers deep. Data was also drawn from quake frequencies of up to 18 hertz, which previous simulations of violent earthquakes could not achieve due to the enormous memory and time requirements needed for high frequency simulations.

The improved simulations, for which the team took home the 2017 Association for Computing Machinery Gordon Bell Prize, may inform new engineering standards for buildings in earthquake-prone zones.

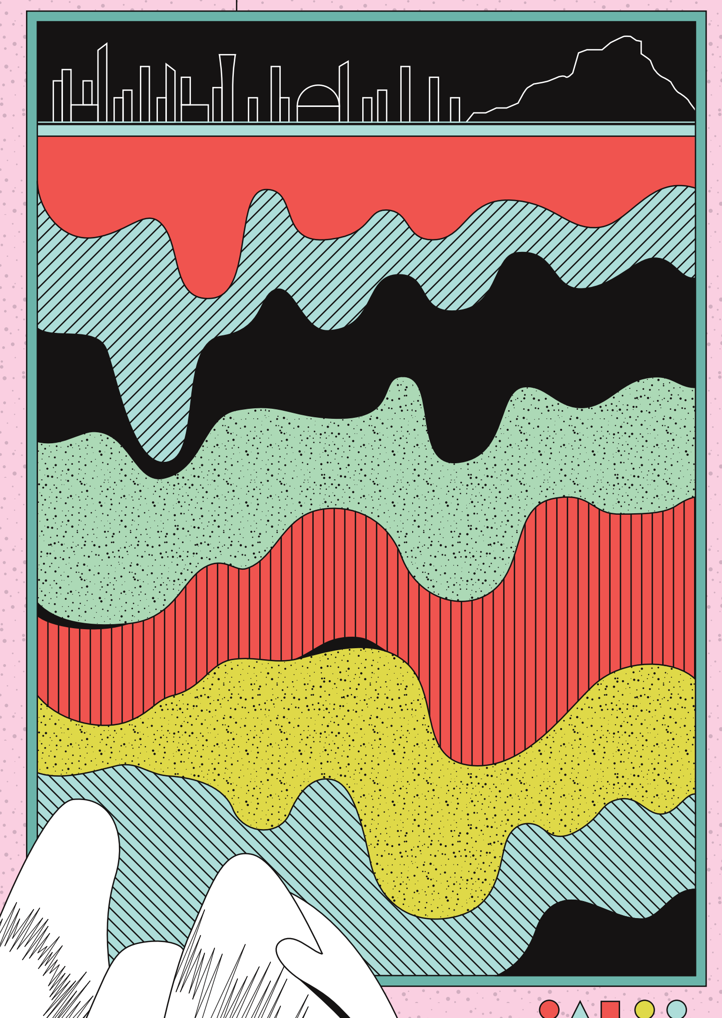

COVER STORY //

SUPPORTING SCIENCE AT ALL SCALES

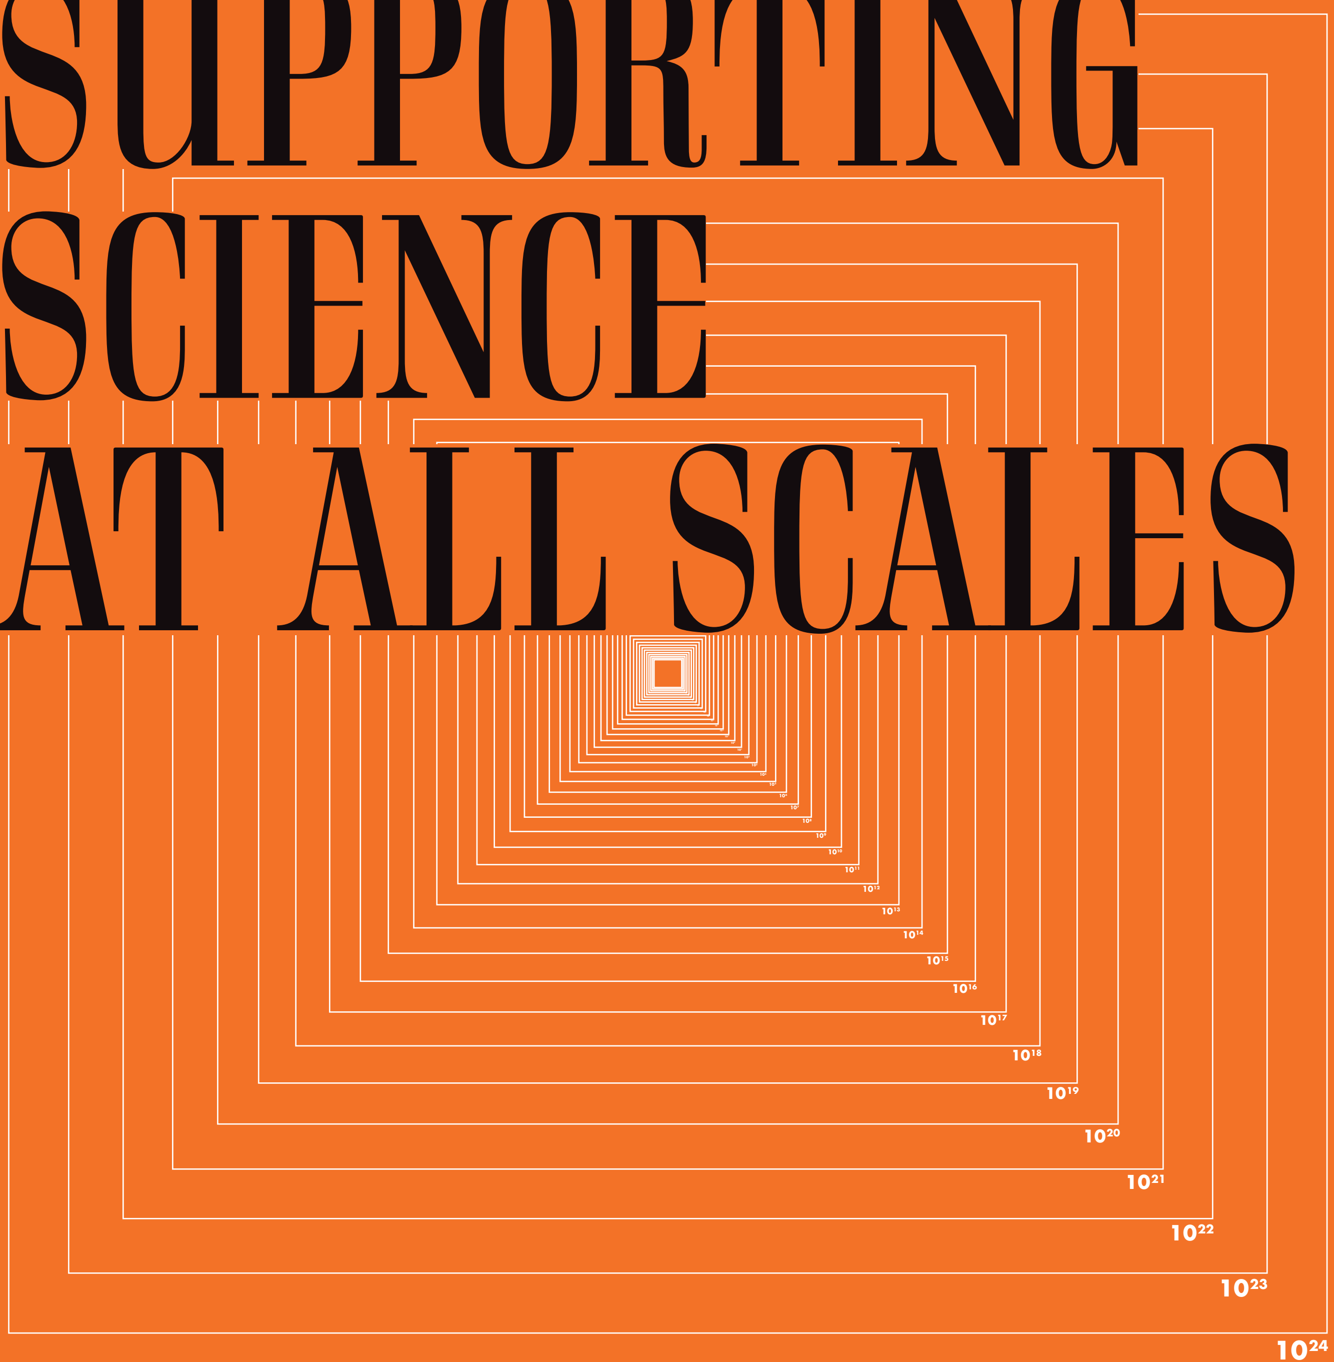

THE SILICON BEHIND THE SCIENCE

Supercomputers are helping scientists investigate phenomena at all scales, from the cosmos to quantum mechanics.

By **Sim Shuzhen**

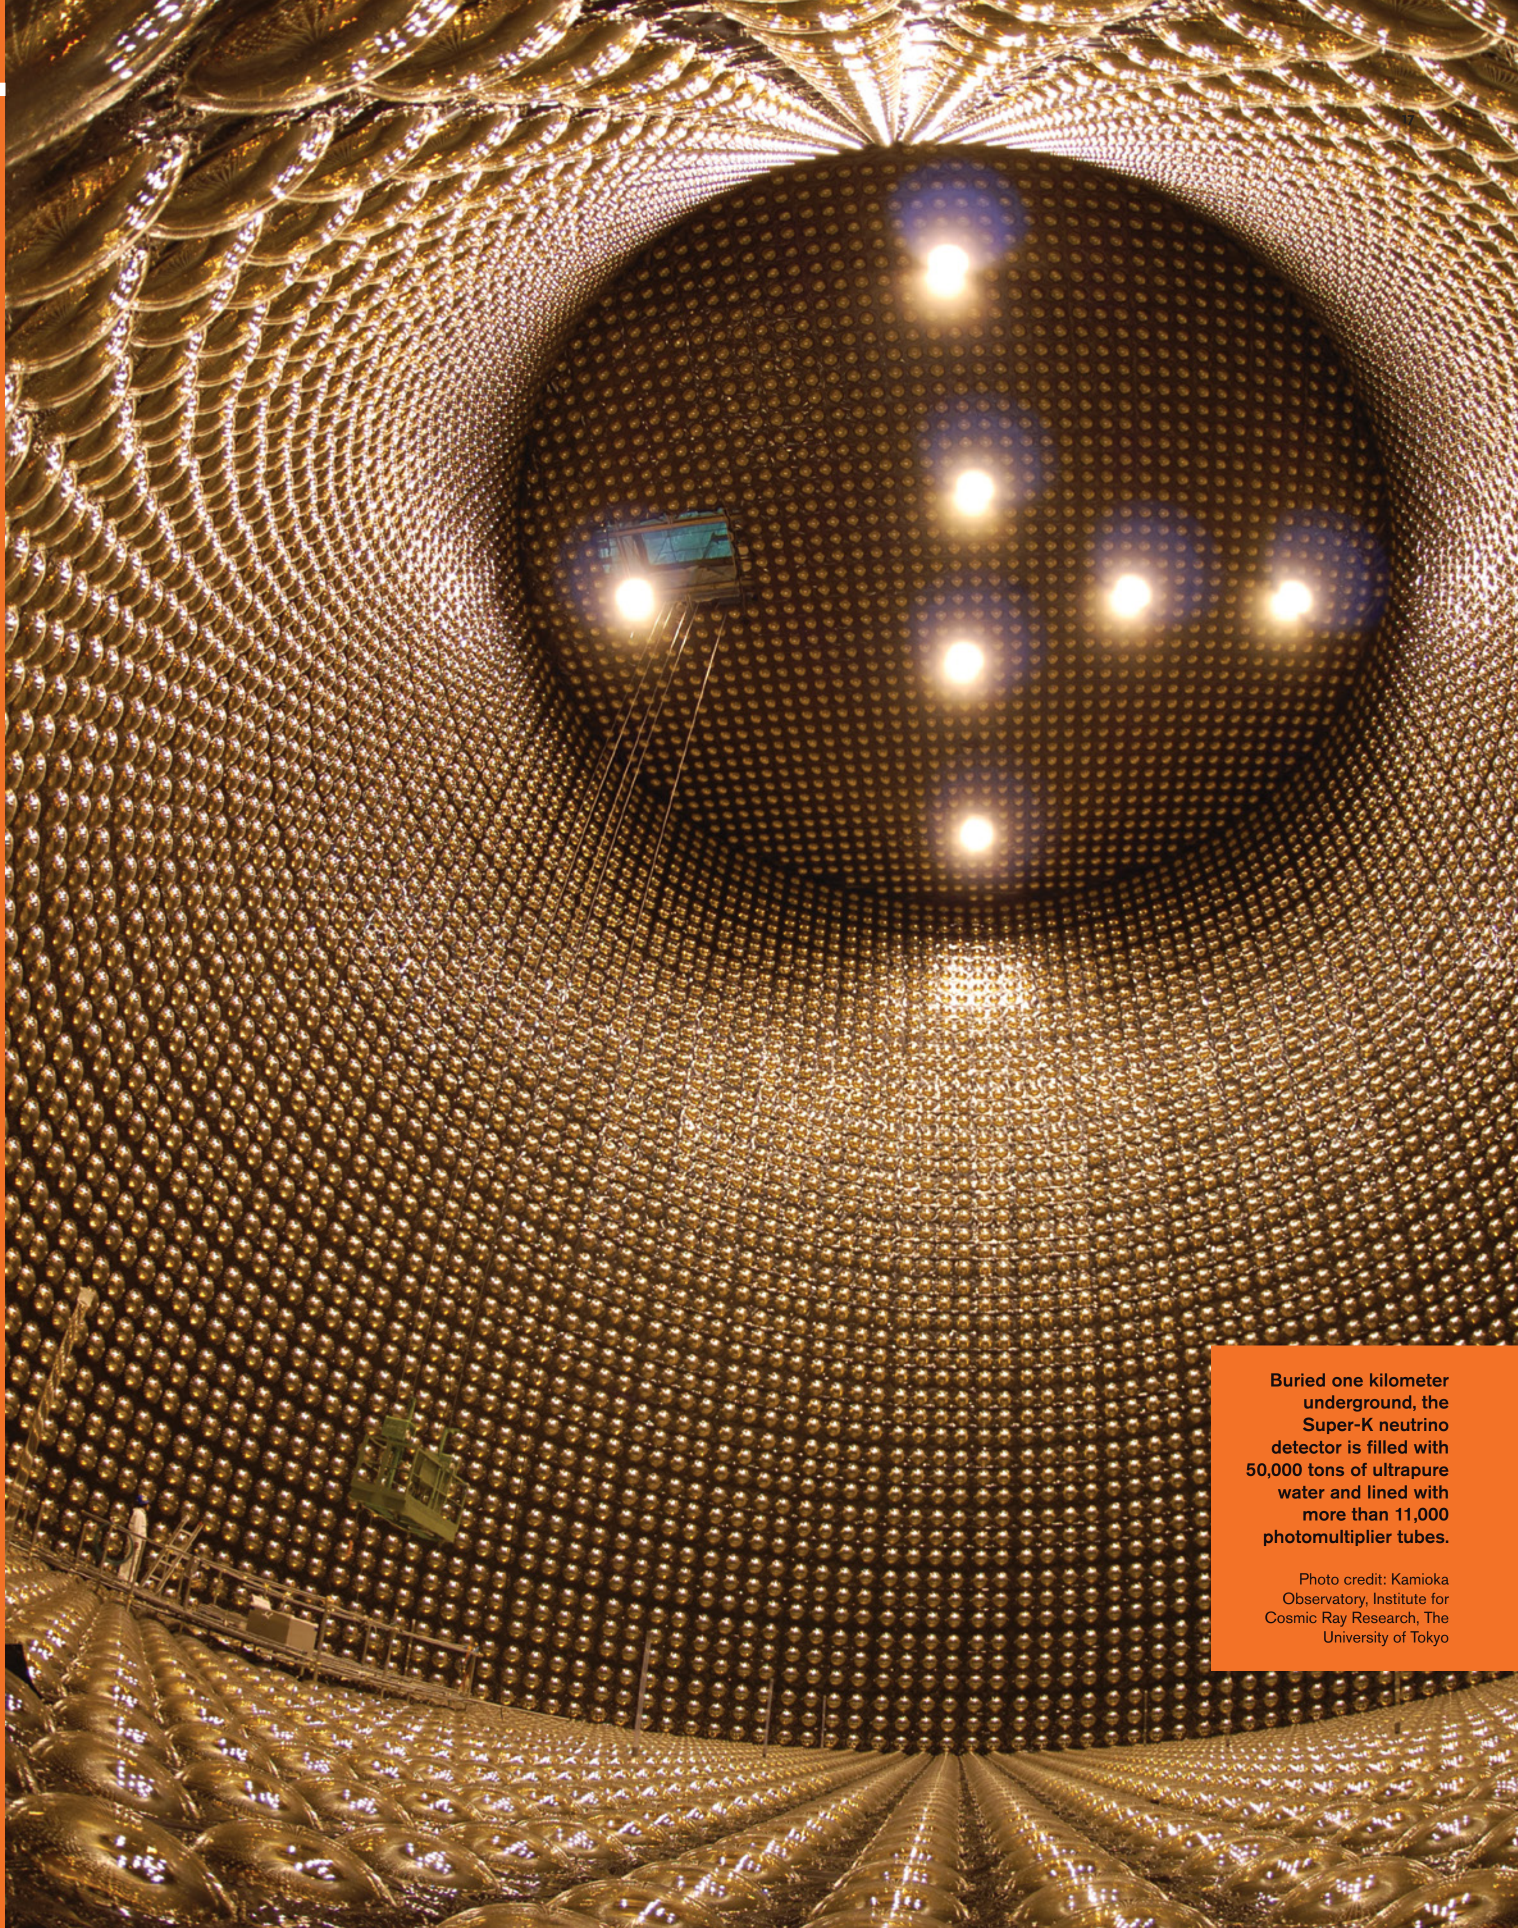

Buried one kilometer underground, the Super-K neutrino detector is filled with 50,000 tons of ultrapure water and lined with more than 11,000 photomultiplier tubes.

Photo credit: Kamioka Observatory, Institute for Cosmic Ray Research, The University of Tokyo

The film begins ordinarily enough, with a couple picnicking on the green lawn of a Chicago park. But then the camera pulls back, the couple getting smaller, to an area ten meters across, then 100 meters, then a kilometer... and so on until it reveals a vista 10^{24} meters across, the estimated size of the observable universe.

Boomeranging back to the couple on the grass, the camera next zooms in on the man's hand, then a layer of skin, then a single cell, then a DNA molecule. The dizzying journey finally ends at a field of view an infinitesimal 10^{-16} meters across; at this magnification, and if such a camera did indeed exist, we would be able to see the quarks within a proton within an atom.

Powers of Ten, a 1977 short film from Charles and Ray Eames—designers better known for their iconic chairs—crams 40 orders of magnitude of spatial scale into a mere nine minutes. Mind-bending as this is, consider that the entire scientific enterprise already spans this spectrum of sizes, whether it is the interstellar considerations of astronomy and cosmology, the planetary proportions of geology and earth systems, the human-sized domains of physiology and medicine, the microscopic worlds of genetics and cell biology, or the subatomic 'spooky' realms of quantum mechanics.

Of course, natural phenomena, as well as the methods and equipment scientists use to understand them, vary widely across different spatial scales. But just about every branch of science is now generating more data than ever before, with an accompanying demand for ever more sophisticated computing power to make sense of it. Supercomputers, then, look set to be a common thread running through science at all scales.

THE RADIO STAR

Probing the vastness of the universe is a project called the Square Kilometer Array (SKA), an effort—also enormous in terms of its own physical size and international scope—to build the world's largest radio telescope, with a collecting area of more than one square kilometer.

Composed of telescope arrays in two locations—Western Australia's Murchison Shire and South Africa's Karoo region—the SKA will let researchers survey the cosmos faster and at higher resolution and sensitivity than ever before. "The SKA is going to be essentially a step change in radio astronomy observatories in terms of capabilities," said Dr. Miles Deegan, a high-performance computing (HPC) and data analytics architecture specialist at the SKA Organization, headquartered at Jodrell Bank Observatory in the UK.

Once the SKA—currently in design phase—starts producing scientifically interesting data in the mid-2020s, researchers will use its superior sky-scanning abilities to study some of the biggest questions in astrophysics, potentially transforming our understanding of the universe, said Deegan.

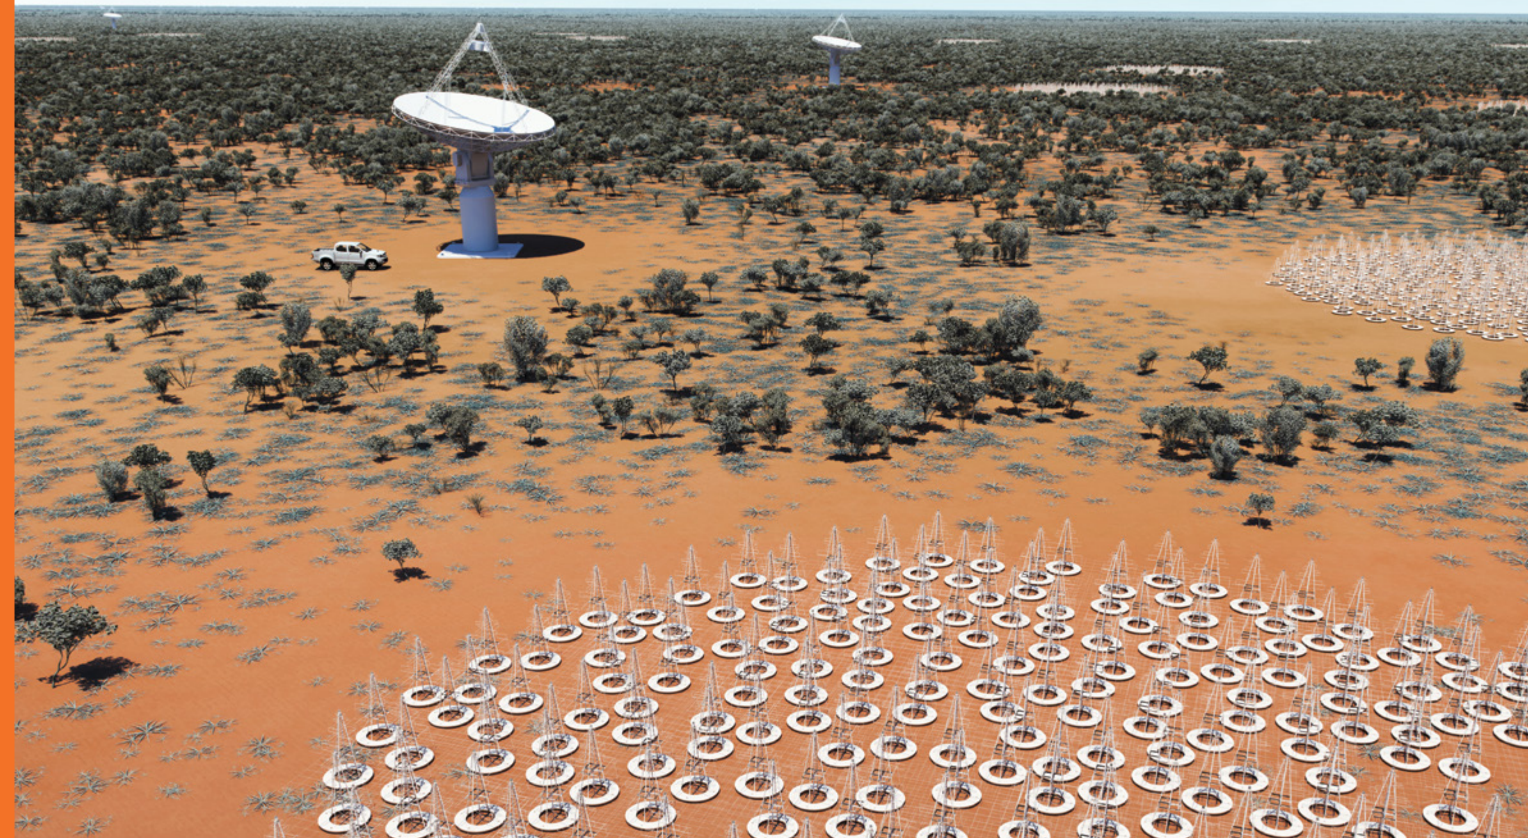

The SKA radio telescope will let researchers survey the cosmos faster and at higher resolution and sensitivity than ever before.

Credit: SKA Organization/Swinburne Astronomy Productions

For instance, by using the SKA to measure the effect of gravity on pulsars that orbit black holes, researchers may be able to detect gravitational waves of different frequencies from those previously discovered through the Nobel Prize-winning efforts of the Laser Interferometer Gravitational-Wave Observatory—work that would further push the boundaries of Einstein's theory of general relativity.

NOT ALL ABOUT THE FLOPS

The computational needs associated with ground-breaking astrophysics research are immense—the first phase of the SKA alone is expected to generate 160 terabytes of raw data per second, approximately five times the global internet traffic in 2015. Processing this barrage of data requires complex supercomputing hardware, software and analytics, which are being designed by the SKA's Science Data Processor Consortium, comprising researchers from institutions in 11 countries.

The two planned supercomputing facilities in Capetown, South Africa and Perth, Australia are likely to have a combined processing power of around 260 petaFLOPS, said Deegan. But acquiring more FLOPS isn't the answer to the SKA's computing challenges, which have more to do with input-output, memory bandwidth and data management, Deegan emphasized.

"It's about the sheer volume of data—capturing that data to start with, being able to ingest it all in time, and managing it through various levels of storage and buffers. That's where the complexity is; the FLOPS are not an afterthought, but it's not all about the FLOPS," he said.

The two supercomputing centers will churn out multiple types of science 'data products,' such as pulsar timing information, said Deegan. But the process doesn't end there—for the data to be fully analyzed and exploited, it then needs to be disseminated to an international network of researchers.

"We need to take about 300 petabytes per annum per telescope and distribute that data to a worldwide network of SKA regional centers... that's another challenge—being able to move vast amounts of data [over] intercontinental distances and distributing to a big network of collaborating institutes," said Deegan. The SKA will likely adopt a model similar to how European particle physics research organisation CERN moves its data internationally via a worldwide grid infrastructure, he added.

DIGGING DEEP FOR NEUTRINOS

Sprawled across two of the world's most remote desert locations, the SKA's antennae fields would be right at home in a science fiction movie. But one doesn't have to look too far to find another research facility rivaling it in scale and strangeness—just one kilometer beneath Mount Ikeno, near the city of Hida, Japan.

There, the Super-Kamiokande (Super-K) detector—a 40-meter-deep steel cylinder filled with 50,000 tons of ultrapure water and lined with more than 11,000 photomultiplier tubes—deals with research at the other end of the spatial scale: it was built to detect neutrinos, one of the most abundant subatomic particles in the universe.

Neutrinos arrive at the detector from various sources, said Professor Masayuki Nakahata, director of the Kamioka Observatory, which hosts Super-K. About a third of the approximately 30 neutrino events Super-K records every day are atmospheric neutrinos, produced when cosmic rays collide with atoms in the atmosphere, while the remaining two thirds are solar neutrinos, generated by fusion reactions in the sun, said Nakahata. These observations have helped researchers make discoveries about the fundamental properties of neutrinos, including Nobel Prize-winning work showing that the particles do indeed have mass.

The subterranean facility, which is operated by an international consortium of some 40 research institutes, is also being upgraded to improve its chances of detecting neutrinos emitted by a third source—supernovae, said Nakahata. Such events, which shed light on the process of star formation, are much rarer and have thus far been detected just once: in 1987, when Super-K's predecessor, Kamiokande, picked up neutrinos from a supernova in the nearby Large Magellanic Cloud. With the upgrade, Super-K is expected to detect neutrinos from supernovae in faraway galaxies, as well as 'supernova relic neutrinos' emitted by stars that exploded in the distant past.

"By observing those supernova neutrinos, we are able to understand the history of the universe—how heavy massive stars exploded and distributed various elements to the universe," said Nakahata.

BLINK AND YOU'LL MISS IT

Despite their abundance, neutrinos are notoriously difficult to detect, since their miniscule mass and lack of electric charge allows them to pass through matter unobstructed, like quantum-scale ninjas. But neutrinos traversing Super-K's huge volume of ultrapure water occasionally do bump into nuclei or electrons, generating charged, high-energy particles that move faster than the speed of light in water.

When this happens, the particles emit a type of radiation called Cherenkov light, which is picked up by the detector's photomultiplier tubes as ring-shaped images. These signals are funneled into Super-K's experiment-analysis system, a Fujitsu HPC setup that includes a data processing system, an 85-server cluster and a high-speed distributed file system located at the underground site.

"We use the computing resources to count the number of rings or observed particles, reconstruct the interaction positions and the energies of each particle, and identify the type of particles from the observed signal," said Associate Professor Yoshinari Hayato, who manages computing at Kamioka Observatory. Further, researchers studying neutrino oscillation—a quantum mechanical phenomenon (discovered in part at Super-K) in which neutrinos switch back and forth between different identities—also need to generate large volumes of simulated data, a computationally intensive process, added Hayato.

Since supernova bursts could happen at any moment, Super-K's top HPC priority is stability, said Hayato—the system cannot afford to blink, or it may miss the chance to record a rare event. "[Supernovae] are known to burst once in decades, and the duration of the burst is just a few seconds... therefore, the number of neutrinos is limited and we need to maximize the opportunity to observe them," he said.

"By observing those supernova neutrinos, we are able to understand the history of the universe—how heavy massive stars exploded and distributed various elements to the universe."

Professor Masayuki Nakahata, director, Kamioka Observatory

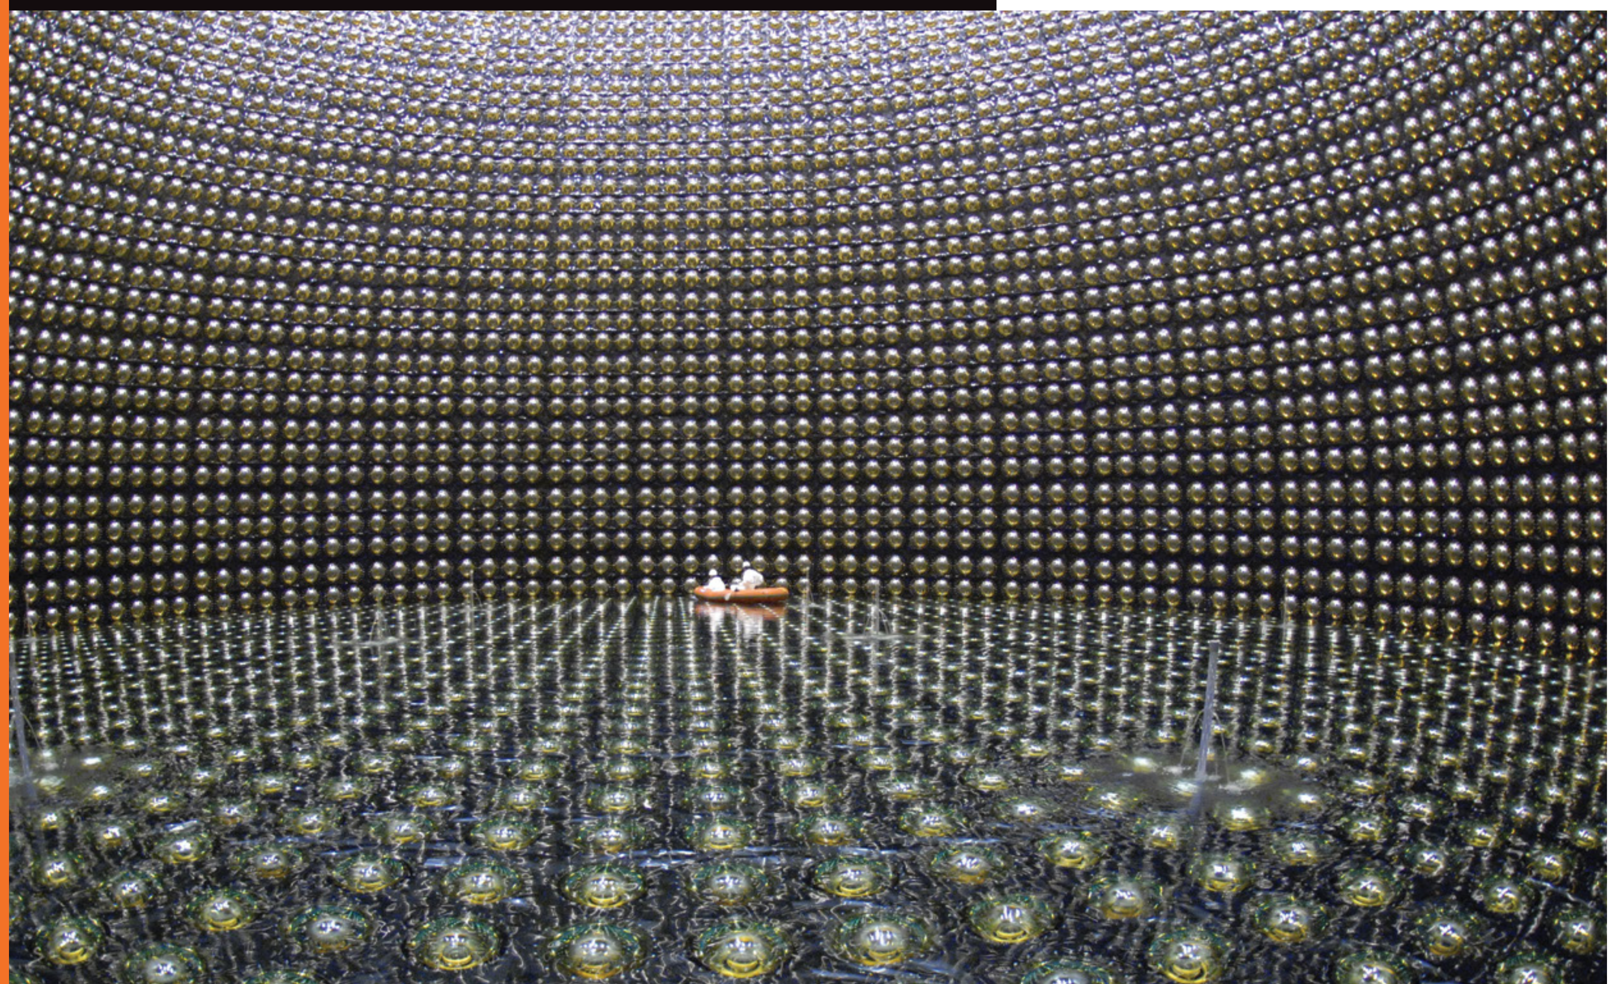

Super-K is being upgraded to improve its chances of detecting neutrinos from supernovae in faraway galaxies.

Photo credit: Kamioka Observatory, Institute for Cosmic Ray Research, The University of Tokyo

SUPERCOMPUTING AND SERENDIPITY

The already-cavernous Super-K will soon be dwarfed by a successor ten times its volume. Kamioka Observatory's planned Hyper-Kamiokande (Hyper-K) neutrino detector, scheduled for completion in the mid-2020s, will allow researchers to address some of physics' most fundamental questions. One of its main goals is to detect a phenomenon in neutrinos known as charge–parity symmetry violation, which could explain why there is more matter than antimatter in the universe.

While official specs are not yet available for Hyper-K's supercomputing infrastructure, Nakahata and Hayato expect it to require orders of magnitude more compute power than Super-K. This will allow it to handle data from the new detector's more than 40,000 photomultiplier tubes (quadruple the number in Super-K), as well as accommodate higher-precision measurements, said Hayato.

Because of the complex nature of the computing infrastructure involved, international scientific collaborations on the scale of the SKA, Super-K and Hyper-K will not only further the science itself, but may also result in new advances in HPC and info-communications technology, said Deegan, referring to the SKA.

"A lot of scientific and technical breakthroughs are often serendipity. But I think the size and complexity of [the SKA] is such that we will have to come up with new ways of doing things. I think [this is] not just on the technology aspects, but also [in terms of] worldwide collaboration... just learning new ways of doing complex science projects," said Deegan.

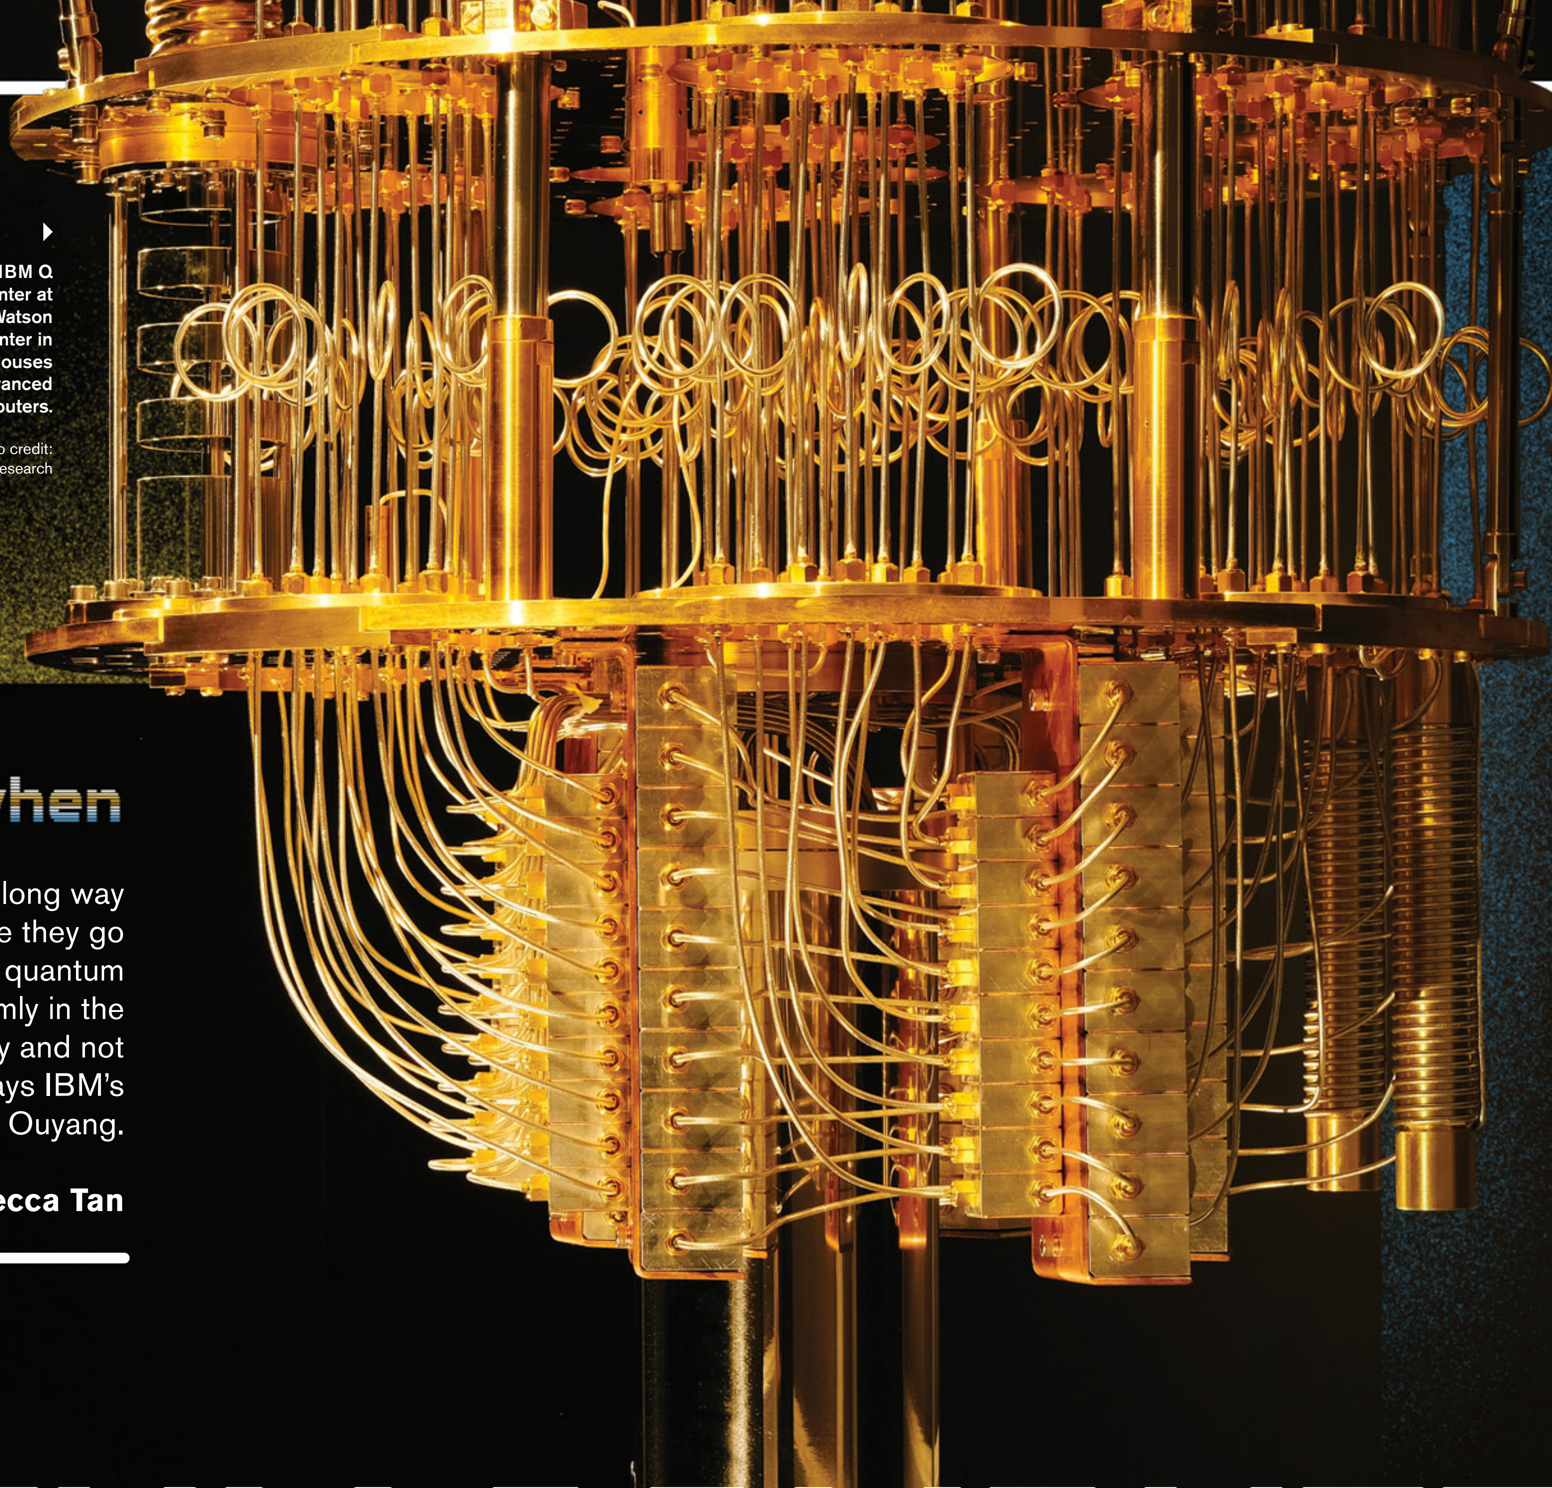

The IBM Q computation center at the Thomas J Watson Research Center in New York, US, houses IBM's most advanced quantum computers.

Photo credit: Connie Zhou/IBM Research

QUANTUM COMPUTERS ON THE HORIZON

Not if, but when

There is still a long way to go before they go mainstream, but quantum computers are firmly in the realm of reality and not science fiction, says IBM's Dr. Christine Ouyang.

By **Rebecca Tan**

Since 1965, Moore's law has served as a roadmap for the chip manufacturing industry, which for decades has kept pace with the original dictum of doubling the number of transistors on a chip every one to two years. However, as we approach the physical limits of miniaturization, it is becoming increasingly clear that a new computing paradigm is needed.

For IBM, getting to 'more than Moore' is likely to involve quantum computers. "We would like to achieve quantum advantage, which refers to the point where quantum computers not only speed up whatever the current classical computers can do, but solve problems that are impossible to solve on classical computers," said Dr. Christine Ouyang, distinguished engineer, IBM Q Network Technical Partnership and Systems Strategies.

"But in order to realize these breakthroughs and make quantum computing both useful and acceptable, we will need to completely re-imagine information processing and the computer that will perform information processing."

A POWERFUL BUT NASCENT TECHNOLOGY

At its simplest level, a classical computer is a collection of bits that are either 0 or 1. The more bits a computer has, the greater the number of possible states it can be in. For example, a single bit has two possible states: either 0 or 1, while two bits have four possible states: 00, 01, 10 or 11.

In a similar way, each quantum bit or qubit has more than one state, analogous to 0 and 1. But unlike classical bits, which are more like a coin placed flat on a table that is either heads or tails, qubits exist in a state known as superposition, which can take on the values of 0, 1, or both at once, like a coin that is spinning and therefore simultaneously heads and tails.

Although a classical computer with two bits has four possible states, it can only be in one of those states at any one point in time. A quantum computer with the same number of qubits, correlated through the phenomenon of quantum entanglement, on the other hand, can simultaneously exist in all four possible states at the same time, making it exponentially more powerful—at least in theory.

At the moment, even the most advanced quantum systems have fewer than a hundred qubits, whereas "business problems might require hundreds, thousands or even millions of qubits," Ouyang said. In a further departure from classical computing, simply adding more qubits is not enough to improve the power of quantum computers.

Another significant challenge standing in the way of quantum computing is how to reduce the error rates of qubits. Part of the solution involves improving the coherence time of qubits—the length of time that researchers can maintain a qubit's quantum state.

To protect them from random interference such as mechanical vibration, electromagnetic waves and temperature fluctuations, a quantum processor's qubits are kept in a dilution refrigerator that is cooled to extremely low temperatures of 10–15 milliKelvin, about a hundred times colder than outer space. Even then, they typically last just a few microseconds—although qubits in IBM Q systems can last as long as 100 microseconds, according to Ouyang—limiting the number of calculations that can be performed.

"If you look at overall system performance, it is actually a very complex metric to assess how powerful a quantum computer is," Ouyang said. "Coherence time is not the only factor; controllability and scalability are all challenging technical problems that need to be solved."

Inside an IBM quantum computer.

Photo credit: Graham Carlow/ IBM Research

BEYOND HARDWARE HURDLES

Aside from the considerable hardware challenges, quantum computing requires an entirely new software stack, starting from the interactions with the actual device at the assembly language level, all the way up to the operating system, midware, applications, and eventually moving to the cloud. "In addition, we also have to develop a very efficient way to map data onto quantum computers and new algorithms that can deliver quantum speed-up for practical applications," Ouyang added.

But the rate-limiting factor might not turn out to be the technology. The soft side of quantum computing matters as well, Ouyang said. "If you think about the cloud or artificial intelligence, adopting these emerging technologies requires organizations to have a certain culture and a set of skills," she said. "It usually takes decades for enterprises to fully embrace new technology."

Recognizing the non-intuitiveness of programming for quantum computers and the need for cultural change, IBM has since 2016 made quantum computing available for free through its cloud-based IBM Q Experience. Over the past two years, IBM's five-qubit and 16-qubit systems have been used by over 100,000 people, including not only scientists and developers but also students.

"Collectively, these people have run more than 6.5 million experiments and published more than 130 research papers. In December 2017, we launched the IBM Q Network and have since grown it into a healthy ecosystem of Fortune 500 companies, start-ups, universities and national research labs," Ouyang said.

"All this tells you is that although it is still in an early stage, quantum computers are here today. It's not science fiction anymore—it is a reality."

THE ROAD AHEAD

Do the advancements in quantum computing then spell the end of traditional high performance computing? For Ouyang, the answer is a categorical no.

First of all, not all problems would benefit from quantum computing. "At least currently, we don't think quantum computers would be good at solving big data problems because they can only take a small number of inputs but explore a large number of permutations simultaneously," Ouyang explained.

"What we have at IBM is a hybrid computing approach. We are building an integrated system that would allow classical computers to tap into quantum computers to solve specific problems."

The ultimate aim, however, is to overcome the aforementioned technical challenges and build what is known as a fault tolerant universal quantum computer, which would provide quadratic or even exponential performance enhancements and achieve true quantum advantage. That lofty goal is likely to take several more decades, Ouyang said.

"But before we get there, we should be able to solve some problems considered unsolvable by classical computers with our current approximate universal computer," she added, pointing to early applications in computational chemistry, financial risk analysis, optimization and quantum machine learning.

In fact, some of the first practical use cases could well come from Asia, where institutions and companies have been working closely with IBM Q researchers. Japan's Keio University, for example, is home to IBM's first commercial hub in Asia, and has helped to bring together partners such as chemical company JSR, Mitsubishi UFJ Financial Group, Mizuho Financial Group and Mitsubishi Chemical to develop quantum applications for business.

"We are working with many industry leaders and universities in the region to apply quantum computing technology, as well as build skills, and at the same time build a market for quantum," Ouyang said.

"It's not science fiction anymore—it is a reality."

QUANTUM COMPUTERS FOR A QUANTUM WORLD

Underlying this drive to take quantum computers mainstream is the fact that nature itself is quantum and best described by a quantum system. As Nobel laureate Richard Feynman famously declared in 1981: "Nature isn't classical, dammit, and if you want to make a simulation of nature, you'd better make it quantum mechanical."

For example, although today's massively parallel computers are already able to run complex chemistry simulations, in reality there are several assumptions and simplifications that have been built into the models.

In pursuit of a quantum picture of nature, IBM Q researchers have used their quantum computing system to simulate beryllium hydride—a relatively small molecule, but the largest one simulated by a quantum computer to date. Their research paves the way for exact representations of larger molecules, a topic of great interest to pharmaceutical companies developing small molecule drugs and researchers searching for exotic new materials.

"The completely new paradigm of quantum computing could someday provide breakthroughs in many other disciplines too, such as optimization of very complex systems and artificial intelligence," Ouyang said. "Our immediate goal is to solve the technical challenges and at the same time work with partners to develop the first wave of applications."

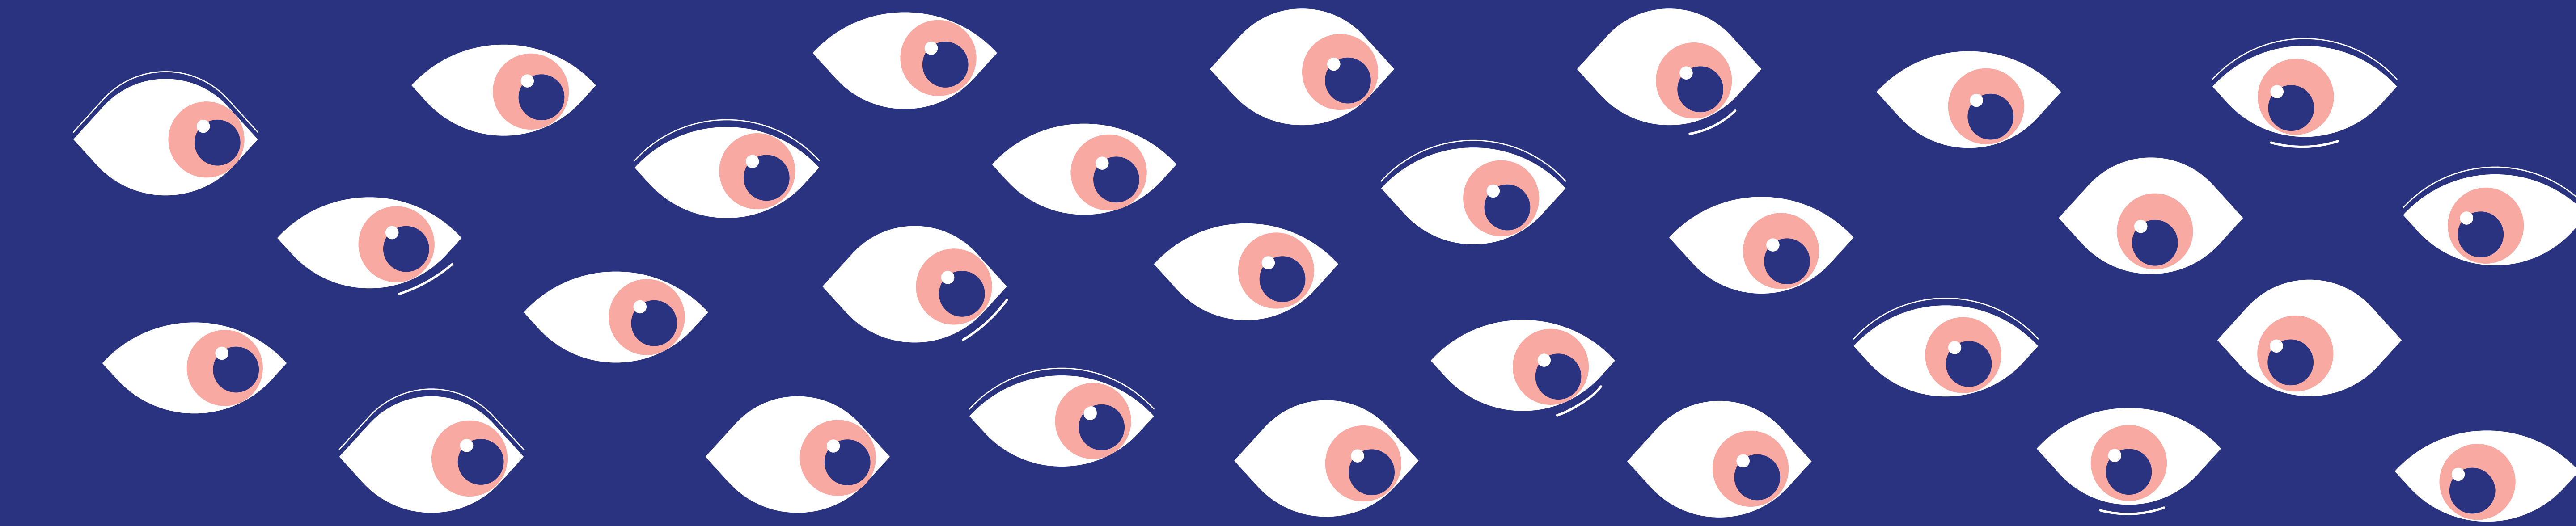

ALL EYES ON THE FUTURE

Bringing sight to machines

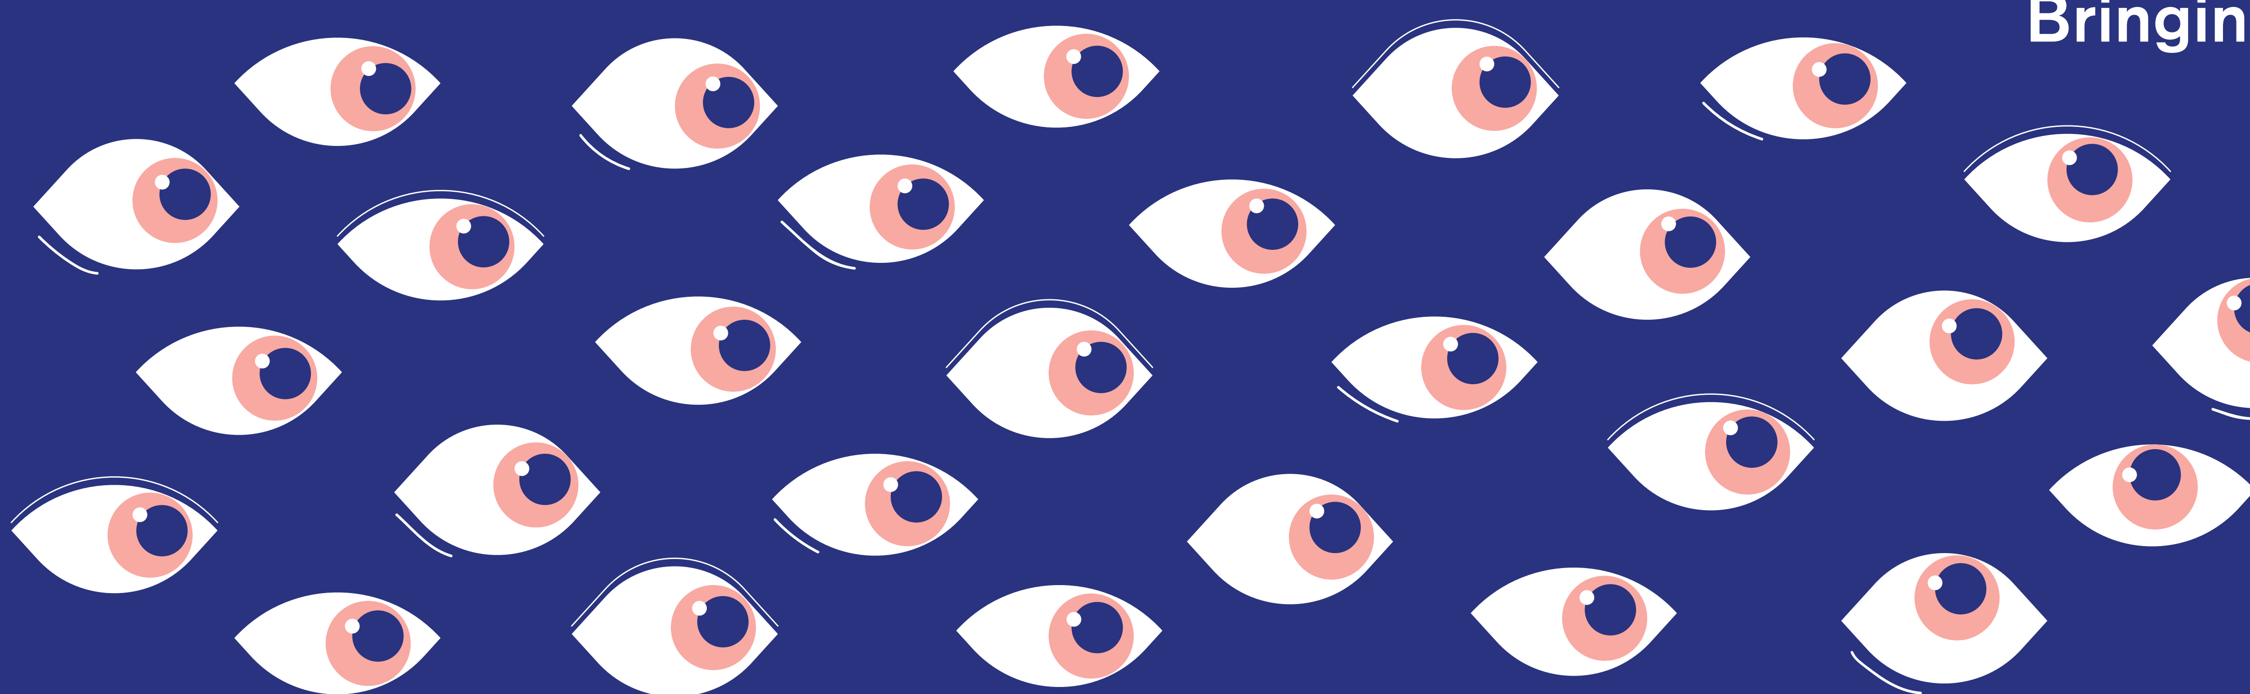

Supercomputers play a key role in artificial intelligence-enabled computer vision, says SenseTime co-founder Lin Dahua.

By **Jeremy Chan**

Illustrations by Lam Oi Keat/*Supercomputing Asia*

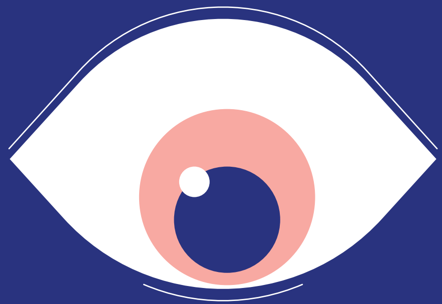

When the first biological eye appeared some 600 million years ago, the rules of survival were irrevocably changed. By being able to detect light and respond to it, organisms that could see were better at seeking out food and avoiding predators. "In the land of the blind, the one-eyed man is king," wrote the 15th century Dutch scholar Erasmus. He was a millennia late in making that observation.

While the sense of sight used to be the domain of living things, humans are now seeking ways to endow our machines with vision. This goes beyond creating lenses for cameras and storage devices for images or videos. What scientists want to do is allow machines to recognize visual information and make sense of it, ideally in real-time. 'Computer vision' is the official term for this technology, and practical applications of it are already emerging. Among the frontrunners in developing and deploying computer vision is Chinese artificial intelligence (AI) company SenseTime.

Founded in October 2014 by Professor Tang Xiaoou of the Chinese University of Hong Kong (CUHK), SenseTime has risen meteorically to become the world's biggest AI startup unicorn, having raised a total of US$1.6 billion in funding over the last four years. Counted among its backers are China's e-commerce giant Alibaba and Singapore's state investment firm Temasek Holdings, which see tremendous potential in SenseTime's core technologies.

"SenseTime is the first to introduce deep learning into the area of computer vision and the first to develop a facial recognition system that surpasses human performance," said Assistant Professor Lin Dahua, director of the CUHK-SenseTime Joint Lab and a co-founder of SenseTime, in an interview with *Supercomputing Asia*. "Our offerings are built on more than 20 years of in-house research and development, and we now have more than 700 partners and customers in China and around the world using our deep learning platform."

A LOOK BENEATH THE SURFACE

By deep learning, Lin is referring to the "tens to hundreds of layers of computation" that allow a machine to recognize an image. In essence, the coded instructions—known as algorithms—tell the computer how to pick out features or patterns in pictures and classify them automatically.

For example, one layer could be detecting the edges of an object in an image, while another could be identifying color or shape. By combining these multiple layers of instructions to create an artificial neural network—so called because it mimics the way the human brain processes information—a machine is able to infer the content of an image, even to the point of captioning or describing it.

But before a computer can make inferences accurately, it needs to be trained, said Lin. Just as a child might learn through reading more books or repeating a series of actions, an artificial neural network is trained with data that is funneled through its layers repeatedly. With each pass, the set of algorithms automatically adjusts its parameters until it eventually becomes very good at its assigned task. This is known as offline learning.

"In the case of facial recognition, you have to run the training procedure over millions of faces, with tens of thousands of iterations, and this is really computationally intensive. Hence, supercomputers play a foundational role in the development of AI-enabled systems," Lin explained.

SUPERCOMPUTERS FOR SIGHT

As a pioneer of deep learning in computer vision, SenseTime naturally has high-performance computing resources at its disposal. With a supercomputing platform comprising more than 6,000 high-performance graphics processing units (GPUs) and a deep learning large-scale training system that allows the simultaneous training of hundreds of GPUs, SenseTime has no trouble crunching huge datasets to optimize its artificial neural networks and achieve accurate facial recognition.

"Typically, the range of computing power we operate with lies between several teraFLOPS to several hundred teraFLOPS," Lin noted. AI training is thus beyond the capacities of the average personal computer. "But once you have optimized the artificial neural network for facial recognition, it can be run in real-time even on a mobile phone."

This is possible because all an optimized algorithm needs to do is convert the image of a face into a string of several hundred numbers (known as a vector), then compare it against a database where every person is already encoded as a vector. "This can be done very efficiently, even for millions of people," said Lin.

Equipped with these capabilities, SenseTime has already made inroads into several industries, Lin added. The Guangzhou Public Security Bureau currently uses SenseTime's intelligent surveillance system to facilitate criminal investigations; according to the company, since 2017, more than 2,000 suspects have been identified and close to 100 cases solved using its technology.

On a lighter note, microblogging website Weibo also relies on SenseTime's facial recognition methods to enhance its user experience and widen its reach. Lin emphasized that "experience working with different vertical sectors is an important aspect that sets SenseTime apart from many competitors."

MORE THAN MEETS THE EYE

But there are areas where the technology still falls short. For instance, under poor lighting, or when a person is "uncooperative" and not looking directly at the camera, the facial recognition system may be less effective.

"Also, when we try to scale up this recognition system to millions of people in a large city, there are bound to be people who may resemble each other, and this makes unique identification tricky. Coupled with the manually intensive process of annotating training datasets, there are still challenges that we have to tackle to take facial recognition to the next level," said Lin.

This is where collaborations with academia play an important role, he commented, hence his appointment as the director of the CUHK-SenseTime Joint Lab. One project that Lin works on is the development of a deep learning framework to identify an individual in a video feed based on just a single portrait image. Rather than consider only visual data, the algorithm factors in temporal information so that a person of interest can be recognized even in environments that are different from where the portrait image was originally taken.

"A lot of research problems require a longer term to explore truly innovative methodologies, and some of that work may be done in the universities," Lin noted, adding that SenseTime's academic collaborations extend beyond China's borders.

In February 2018, SenseTime partnered with the Massachusetts Institute of Technology, US, to leverage the latter's strengths in AI research. Shortly after, it also signed memoranda of understanding with three Singapore-based organizations—the National Supercomputing Centre Singapore, Nanyang Technological University and telecommunications provider Singtel—as part of its international expansion plans.

Having helped computers open their eyes, SenseTime looks set to take the technology to the masses. Like the first animals with sight, machines may have a very different vision for the future, and it remains to be seen where humans fit into that picture.

DEEP LEARNING X COMPUTER VISION

Augmented reality
Human-computer interaction tools

Autonomous driving
Sensing of road conditions

Deep learning chips
Custom hardware for high-performance computation

Financial services
Online remote identity verification

Healthcare
Analysis of medical images

Mobile phones
Apps

Smart cities
Surveillance

"SUPERCOMPUTERS PLAY A FOUNDATIONAL ROLE IN THE DEVELOPMENT OF AI-ENABLED SYSTEMS."

—Assistant Professor Lin Dahua, director of the CUHK-SenseTime Joint Lab and co-founder of SenseTime

SUPERCOMPUTING SHOULDN'T BE ROCKET SCIENCE

Putting computing power in the hands of scientists

Given the rapid pace of data generation, supercomputers are fast becoming an indispensable tool for research—now the challenge is to help scientists use them.

By **Jeremy Chan**

From massive systems to minuscule structures, scientists are collecting data at every imaginable scale: satellites and sensors constantly survey our planet and the cosmos, while scientists in laboratories around the world scrutinize the movements of individual atoms and molecules.

But data in its raw form is like unprocessed mineral ore—ugly and not particularly useful. The fact that there is so much of it only compounds the problem—sifting through the noise to derive insights from data is not only expensive and time-consuming, but also almost impossible to do manually. As the rate of data generation outstrips the pace of analysis, scientists are finding supercomputers increasingly necessary to conduct meaningful research.

"For example, modeling the Big Bang and forecasting the weather are very computationally intensive. Also, in biology, the number-crunching capabilities of supercomputers are needed for modeling molecular dynamics," said Dr. Andreas Wilm, team leader of the bioinformatics core at the Genome Institute of Singapore (GIS). "There are plenty of research areas that require high performance computing (HPC)."

HPC AS A PUBLIC UTILITY

Useful as supercomputers may be, not everyone has access to them. Here, access may be defined in two different ways—availability of supercomputing resources, and technical ability to use supercomputing productively.

Most of the world's HPC prowess is currently concentrated in the US, China and Japan, which means that one does not simply plug in to a petaFLOP supercomputer and run an experiment. This is set to change as cloud computing providers such as Microsoft Azure increasingly offer HPC-like resources on demand, said Wilm.

"Rather than own supercomputing hardware themselves, researchers can tap on systems like AWS' CfnCluster or Microsoft's Cycle Computing to spin up their very own 'supercomputer' in the cloud from their laptops," he added.

Mr. Alex Nodeland, co-founder and CEO of Archanan, a company that develops cloud-based development environments for HPC software, thinks that in the future, "the majority of the world's HPC resources will be accessible over a cloud-like interface." This looks set to usher in an era of democratized supercomputing, where geography and cost no longer prevent scientists from performing complex and computationally intensive analyses.

FLOPS IS NOT WHERE IT STOPS

Hardware woes aside, the processes and practices surrounding the use of supercomputers also pose barriers to the widespread adoption of supercomputing for research.

"For complex analytics such as those used in computational genomics, researchers may have to execute many different programs on different slices of data in a specific order, and with certain dependencies between jobs. Managing this orchestration can become an overwhelming task," explained Wilm.

Nodeland also highlighted this steep learning curve for researchers. "HPC users come from diverse backgrounds and are not always well equipped with knowledge of computer architecture and networking. Combined with very diverse user interfaces and complex computing environments, it can be challenging for scientists to take advantage of supercomputers for their research."

SUPERCOMPUTING SIMPLIFIED

Fortunately for researchers, the HPC community is working to improve the usability of supercomputers. To overcome obstacles in orchestration, for instance, Wilm recommended the use of workflow management software such as Snakemake and Nextflow, which simplify the development, execution and scaling of a computational pipeline.

"With a given workflow recipe, users—in theory—only have to invoke one command, which then automatically orchestrates hundreds of others," he said, alluding to a domino effect where an initial trigger drives all downstream activities. "At GIS, we use Nextflow to orchestrate the analysis of thousands of genomic samples on the National Supercomputing Centre's ASPIRE 1 supercomputer."

Meanwhile, Nodeland is looking to go beyond workflow management and enable researchers to custom-build their own HPC software. Yet, software created on traditional desktops may not run properly when installed on supercomputers due to differences in core count and coding standards. Because of this, about one third of a supercomputer's capacity typically goes towards development and debugging, hogging HPC resources that could otherwise have been used for running productive computations, he said.

"We need to bridge the gap between the creation of large-scale parallel codes on conventional workstations and the deployment of software on multi-million-dollar supercomputers," he added.

SIMULATION STATION

Nodeland's solution involves getting researchers to offload their software engineering processes to the cloud. His company, Archanan, has built a platform consisting of tools such as debuggers, profilers and memory map analytics, which allows researchers to develop, test and deploy HPC applications on cloud-based infrastructure that emulates a supercomputer.

With this platform, researchers can have greater confidence that whatever software they code on the cloud will function as intended when ported over to an actual supercomputer, since the development (cloud) and deployment (supercomputer) conditions are now largely similar.

Nodeland told *Supercomputing Asia* that early adopters of Archanan's platform are already seeing improvements in development life cycles for supercomputing software pertaining to materials science, bioengineering, physics, economics, data mining and artificial intelligence, among others. "By removing the burden of dealing with system architecture, Archanan allows HPC users to focus on their research, leading to better results in less time," he concluded.

EMBRACING AI IN SINGAPORE

How the little red dot is getting AI-ready

Efficient, accessible and flexible supercomputing resources will play a key role in helping Singapore harness AI technology, says AI Singapore's Professor Leong Tze Yun.

By **Brenda Lau**

Today, there's no escaping the fact that artificial intelligence (AI) has crept into many aspects of our lives, making communications easier, workflows more effective and research analyses much faster.

Recognizing the productivity gains to be had from the power and pervasiveness of AI, Singapore intends to embrace the technology as part of its Smart Nation initiative. But are businesses and consumers ready to hop onto this AI-adopting journey?

To find out, *Supercomputing Asia* talked to Professor Leong Tze Yun, director of AI technology at AI Singapore (AISG), a national initiative set up to coordinate the city-state's capabilities in the technology.

What is the state of AI activity in Singapore?

Prof Leong Tze Yun: Singapore is no stranger to the adoption of AI and takes pride in having a very active research and development community in AI and its related disciplines. The activity is fueled by highly regarded research work done by universities and research institutions.

A clear sign of adoption is also seen through many successful startup companies that develop or deploy AI technologies. Some of these have even originated from AI research and development in Singapore, for example, companies like [visual search and image recognition startup] ViSenze and [workforce optimization startup] FriarTuck.

What challenges do businesses and end-users face in getting technology-ready for AI?

LTY: The most common struggle is the uncertainty of where and how to start with AI. Many people also have doubts and misunderstandings and are often influenced by popular literature and myths. In fact, AI technologies have already been changing our daily lives through the way we work and communicate with others—such as predictive analytics in medical diagnoses and natural language processing in chatbots.

I believe the key to addressing this challenge is in talent development, recruitment and deployment. Apart from public education, a more concerted effort needs to be made in the formal education and training of AI scientists and engineers at the tertiary level, as well as in sectors such as finance, healthcare, education and logistics.

A new model of lifelong learning for the workforce has emerged. For example, industry and institutes of higher learning are working closely together to retrain and upskill employees with specialized skills, and to provide students with more internship opportunities. The government is also providing incentives and establishing policies and guidelines for implementing lifelong learning programs and specialized programs to reach out to targeted sectors, the schools and the public.

How is AISG boosting Singapore's AI capabilities?

LTY: As AISG is a government-driven national program, our goal aligns with the government's ambitions for AI, which is to catalyze, synergize and boost Singapore's AI capabilities to power our future digital economy.

AISG will bring together all Singapore-based publicly funded research institutions and the vibrant ecosystem of AI start-ups and companies developing AI products, to perform fundamental and use-inspired research, grow the knowledge, create the tools and develop the talent to power Singapore's AI efforts.

We support AI innovation by co-funding proof-of-value projects in our 100 Experiments (100E) program. To accelerate adoption of AI by industry, the program pairs companies that have a significant industry problem with AI researchers in Singapore to work on their problem statement.

We also support a new approach to research and education in the AI Grand Challenge program, which addresses highly complex issues with significant social and economic impacts in Singapore and beyond. We aim to provide a translation framework and test bed for multidisciplinary teams to work together to devise innovative solutions that can be put into practical use.

What role do supercomputers play in Singapore's AI strategy?

LTY: I would define the future of 'supercomputing' for AI as a highly efficient, highly accessible and highly flexible computing and information framework.

In addition to the clusters of high performance processors and data centers, future computing and information platforms for AI research, development, translation and innovation should support rapid prototyping, creative experimentation and effective translation to full-scale production in a fast, secure and adaptable manner.

The platforms and frameworks should facilitate collection, communication, analysis and generation of information from human experts and users as well as from edge and mobile devices. They should also support automated and interactive prototyping, real-life experiments and large-scale implementation.

In other words, the underlying computing and information infrastructure should support the whole process of idea generation, knowledge discovery, real-time test-bedding and scaling for deployment.

How can AISG help researchers access supercomputing resources?

LTY: We are working closely with the National Supercomputing Centre Singapore to support the work done by researchers and companies participating in our programs. We also provide support for research into the development and deployment of AI technologies, applications and systems.

In fact, we are building an in-house, high-performance computing platform to support a range of in-house and funded research and education activities. With our partners, we are also assessing the feasibility of establishing a secure, privacy-preserving and protected access facility to provide relevant real-life data for the teams working on the AI Grand Challenges.

INSPIRING MINDS AT SC18

Highlights from Dallas, Texas

SC18, held from November 11–16, 2018 in Dallas, Texas, marked the 30th anniversary of the international conference for high performance computing, networking, storage and analysis. Themed 'HPC Inspires,' this year's conference focused on the power and promise of high performance computing to solve the world's most difficult challenges.

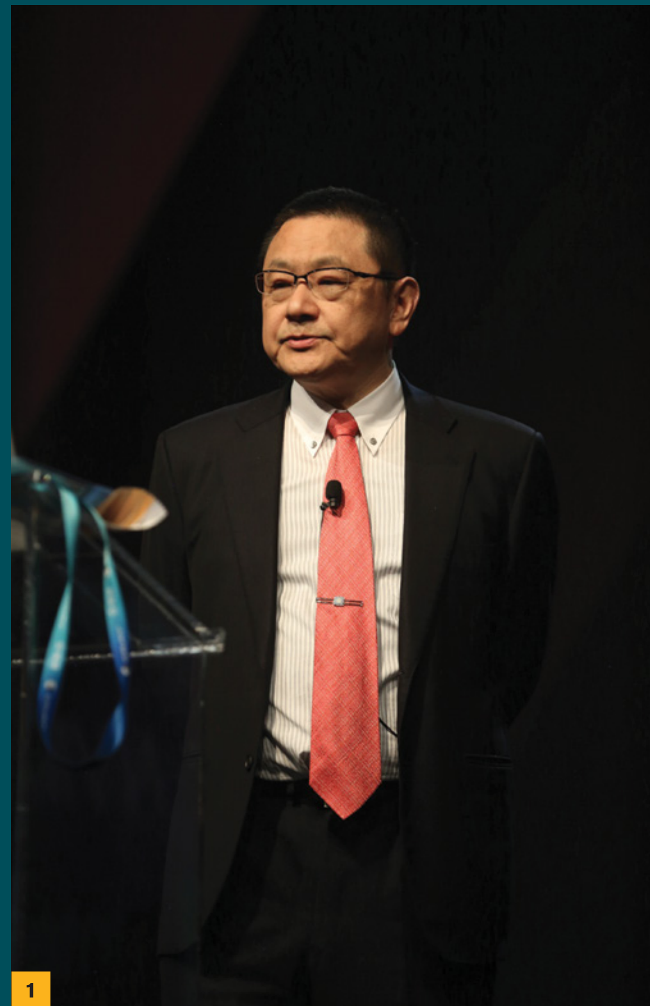

1

2

3

4

5

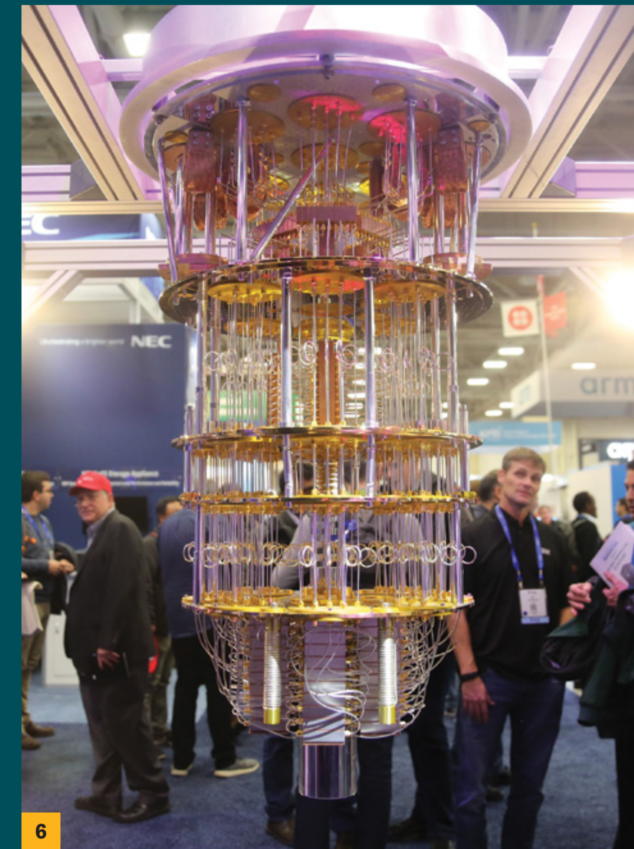

6

7

8

1. Dr. Satoshi Sekiguchi, vice president of Japan's National Institute of Advanced Industrial Science and Technology (AIST), discusses the ABCI, AIST's AI supercomputer.
2. Students from Nanyang Technological University, Singapore, achieved the highest LINPACK benchmark in the Student Cluster Competition.
3. Thinking big at SC18.
4. Professor Erik Brynjolfsson of the Massachusetts Institute of Technology delivered the SC18 keynote, speaking on how economies, societies and organizations will be affected by the machine age.
5. Professor Qian Depei of Sun Yat-sen University and Beihang University, China, giving an update on the country's efforts to develop exascale systems.
6. Just hanging out with a quantum computer prototype at the exhibition hall.
7. Dr. Marguerite Nyhan, research scientist at United Nations Global Pulse, speaking at the SC18 plenary. The session focused on how HPC and AI can be harnessed to solve humanity's grand challenges.
8. Crowds thronged the exhibition hall to examine the latest HPC hardware, software and products.

Photo credits: SC Photography

Business Bytes

D-WAVE QUANTUM COMPUTER SIMULATES TOPOLOGICAL STATE OF MATTER

Canadian quantum computing company D-Wave Systems has demonstrated a topological phase transition using its 2048-qubit annealing quantum computer. This complex quantum simulation of materials, published in the journal *Nature*, is a major step toward reducing the need for time-consuming and expensive physical research and development.

In Nobel Prize-winning work done in the 1970s, theoretical physicists predicted a new state of matter characterized by non-trivial topological properties. D-Wave researchers demonstrated this phenomenon by programming the company's quantum computer to form a two-dimensional frustrated lattice of artificial spins. The observed topological properties in the simulated system cannot exist without quantum effects and closely agree with theoretical predictions.

"This is a significant step toward reaching the goal of quantum simulation, enabling the study of material properties before making them in the lab, a process that today can be very costly and time consuming," said Dr. Mohammad Amin, chief scientist at D-Wave.

CRAY BAGS TWO SUPERCOMPUTER CONTRACTS IN ASIA

Research organizations in Japan and South Korea have purchased two XC50 supercomputers from US supercomputer manufacturer Cray.

The first machine, an 815-teraFLOPS system, has been installed at the Railway Technical Research Institute (RTRI), the research and development arm of Japan Railways Group, and will be used for research in new technologies and materials, simulation techniques and evaluation methods. "By fully utilizing this powerful system, we will be able to generate game-changing insights that will fuel one of the world's most advanced railway systems," said Dr. Norimichi Kumagai, president of RTRI.

The second machine has been purchased by the Institute for Basic Science (IBS), South Korea, for US$9 million. The 1.45-petaFLOPS system will help IBS expand its research in climate physics. "With Cray's advanced system design and expertise in high performance computing environments, our researchers are well equipped to address some of the greatest challenges in oceanography, atmospheric science, glaciology and other areas," said Professor Axel Timmermann, director of the IBS Center for Climate Physics.

FUJITSU DISCLOSES DETAILS ON POST-K SUPERCOMPUTER PROCESSOR

Fujitsu has revealed specifications for the processor that will power Japan's first exascale supercomputer, also known as Post-K. The machine, commissioned by RIKEN, is expected to run 100 times faster than the 10.5-petaFLOPS K Computer, also housed at RIKEN.

Speaking at the August 2018 Hot Chips conference in California, US, Mr. Toshio Yoshida, processor architect at Fujitsu, announced that the Post-K processor, an A64FX ARM CPU comprising 8.8 billion transistors manufactured with 7nm process technology, will reach a peak double-precision (64-bit) performance of over 2.7 teraFLOPS. Fujitsu has also boosted the A64FX's performance for 16-bit and 8-bit operations, making it well suited for deep learning and big data applications.

The A64FX will be the first CPU to implement ARM's scalable vector extension, an instruction set designed specifically for high performance computing; Fujitsu and RIKEN will also co-develop software for both the A64FX and for the system itself. Post-K is scheduled to begin operations in 2021.

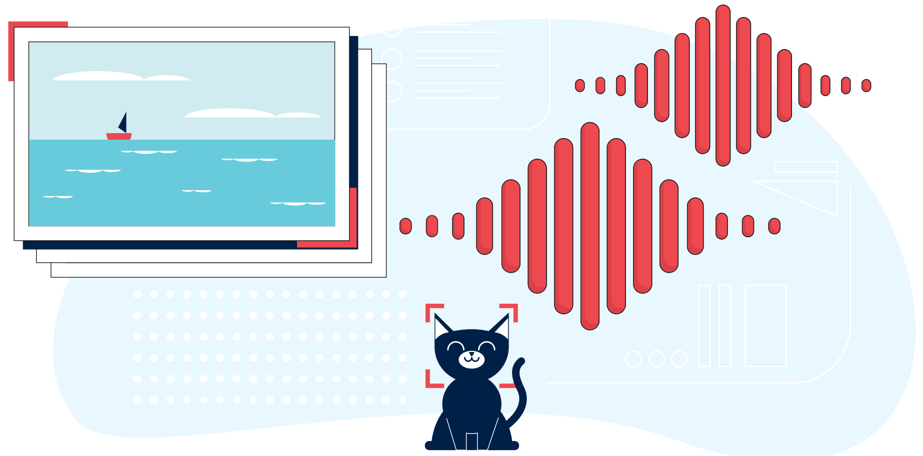

BAIDU LAUNCHES 'NO CODE' PLATFORM FOR DEEP LEARNING

Chinese search giant Baidu has announced the release of EZDL, a no-code platform for building custom machine learning models. Targeted at small- and medium-sized businesses which may lack dedicated coding expertise and large datasets, EZDL features a drag-and-drop interface and can be used for image classification, object detection and sound classification.

"Even if you have had no exposure to programming, you can quickly build models on this platform with zero barriers. EZDL can help companies with limited AI experts and computing resources to quickly and efficiently complete deep learning training and deployment with only a small amount of data," said Mr. Xie Yongkang, EZDL's tech lead.

EZDL is part of Baidu's strategy to develop a comprehensive AI ecosystem; the company's Baidu Brain encompasses a full tech stack from chips to deep learning frameworks to platforms, said Mr. You Youping, Baidu's general manager of AI ecosystem division. "We seek to create a true ecosystem for AI, democratizing access to AI capabilities."

Super Snapshot

MOORE'S LAW: BEYOND THE WALL

After five decades of exponential increases in computing power, Moore's law is starting to hit a wall. For the tech industry, it's time to innovate or stagnate.

IN 1965, Intel co-founder Gordon Moore observed that transistors were shrinking so fast that twice as many every year—later revised to every two years—could fit onto a computer chip.

This insight gave us Moore's law, which predicts that computing power will grow exponentially, doubling every two years.

MOORE'S LAW: TRANSISTORS PER MICROPROCESSOR

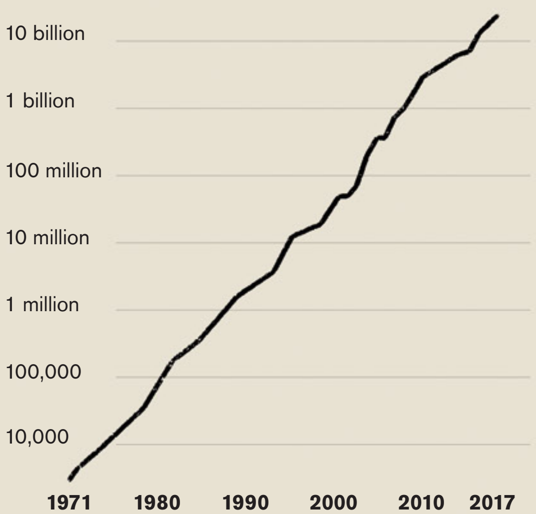

Data available at https://ourworldindata.org/technological-progress (CC licence)

Put into practice, the predictions of Moore's law have for over 50 years enabled ever smaller and faster computers and mobile devices.

But physical limitations on chip manufacturing mean that exponential growth cannot continue forever. To keep up with the demand for computing power, we need to go beyond Moore's law—by developing new ways of doing computing.

POST-MOORE COMPUTING EFFORTS INCLUDE THE DEVELOPMENT OF

- New algorithms to optimize chip performance
- New materials like graphene and silicon-germanium to further reduce transistor size and electrical resistance
- New chip architectures, such as graphics processing units (GPUs) and field-programmable gate arrays (FPGAs)
- Quantum computers, which go beyond the binary encoding of today's machines