AsianScientist

Issue 06
July 2019

SUPERCOMPUTING ASIA

1,000,000,000,000,000,000

THE RACE TO EXASCALE

GETTING ARTIFICIAL INTELLIGENCE AND BIG DATA UP TO SPEED

+

IT'S A MATERIAL WORLD

A QUANTUM LEAP INTO THE FUTURE

MOVE IT!

MCI (P) 037/04/2019

CONTENTS

Issue 06
July 2019

COVER STORY

FEATURES

EDITOR'S NOTE

In this issue, we celebrate yet another successful edition of the annual SupercomputingAsia conference, which saw more than 700 research and industry delegates from around the world flock to Singapore on March 12-14, 2019 (*Scaling New Heights At SCA19*, p. 32). We take a look at India's National Supercomputing Mission, a seven-year project with a budget of US$640 million to build a network of 73 high performance computing centers—three of them petascale—linking research institutions across India (*India's Need For Speed*, p. 38).

Materials scientists have traditionally been one of the biggest users of supercomputers and continue to harness the power of high performance computing to make new breakthroughs. Check out *It's A Material World* (p. 10), to find out how they are designing exciting new materials from radiation-tolerant ceramics to super hydrophobic polymers for cleaning up oil spills. In *A Quantum Leap Into The Future* (p. 24), find out whether advances in quantum key distribution mean that we will soon enjoy an 'unhackable' quantum internet.

Finally, we turn our attention to the most exciting sport in supercomputing—the race to build the world's first exascale supercomputer. Japan is en route to launching its first exascale computer—Post-K is set for launch between 2021 and 2022, with Fujitsu announcing that its design has already been completed. The US' exascale computer, Aurora, is scheduled for delivery by Intel in 2021. How will China stack up in this race to exascale? Find out in our interview with Professor Lu Yutong, director of the National Supercomputing Center in Guangzhou, China (*The Race To Exascale*, p. 18).

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

/asianscientist @asianscientist
Asian Scientist Magazine asianscientist

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

DIGITAL EDITOR
Dr. Jeremy Chan

CONTRIBUTORS
Barry Lo
Eugene Low
Yvonne Gao
Li Lidao

SENIOR DESIGNERS
Lam Oi Keat
Muhamad Azlin

SALES & MARKETING
Samantha Yeap
Grace Phua

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

JAPANESE-US TEAM CREATES FIRST BILLION-ATOM SIMULATION OF DNA

An international research team has created the largest simulation to date of an entire gene locus of DNA, a feat that required one billion atoms to model.

DNA is the blueprint for all living things. With enough DNA inside us to wrap around the earth 2.5 million times, this string-like molecule has to be compacted and wound up in a network of tiny spool-like structures called chromatin.

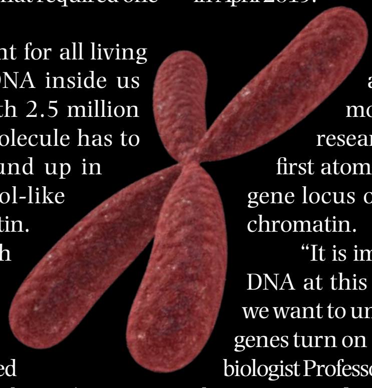

The team, which included researchers from the RIKEN Center for Computational Science in Japan, performed the simulation on Los Alamos' Trinity supercomputer—the sixth fastest in the world with a peak performance exceeding 41 petaFLOPS. Their findings were published in the *Journal of Computational Chemistry* in April 2019.

Using both experimental data and novel computer modeling algorithms, the researchers constructed the first atomistic model of an entire gene locus of GATA4 packaged in chromatin.

"It is important to understand DNA at this level of detail because we want to understand precisely how genes turn on and off," said structural biologist Professor Karissa Sanbonmatsu, the corresponding author of the study. "Knowing how this happens could unlock the secrets to how many diseases occur."

US TO DISCOVER NEW MATERIALS USING SUPERCOMPUTERS

The US Department of Energy (DOE) has announced US$24 million in research funding to develop sophisticated software for the computer-based design of novel materials with applications in energy, electronics, and other areas.

"DOE national laboratories are now home to the two fastest supercomputers in the world, as well as five out of the fastest ten," said the US secretary of energy Mr. Rick Perry. "One major benefit of these advanced systems will be our capability to accelerate the process of innovation in materials through computational design."

US national laboratories, universities, and non-profit organizations will be eligible to compete for the four-year award, which will be selected on the basis of peer review. The Department's Office of Science, which is funding the effort, envisions large teams aiming to develop entirely new software applications, and smaller teams working on research that can add software functionality.

Computational design of new materials is a comparatively recent development, a product of both dramatic advances in computing capabilities and improvements in the theory and understanding of material properties.

The focus of this effort will be on quantum materials, materials related to quantum information science, topological materials, wide-bandgap semiconductors, transition metal oxides, and magnetic, superconducting, photovoltaic, optoelectronic and thermoelectric materials.

UK'S 'MICHAEL' SUPERCOMPUTER TO ACCELERATE BATTERIES RESEARCH

A new supercomputer specifically dedicated to speeding up battery research in the UK has been installed at University College London.

Named Michael, after the UK's most famous battery scientist, Michael Faraday, the supercomputer will reach 265 teraFLOPS at peak performance.

Purchased by the Faraday Institution at £1.6 million (~US$2 million) through the UK Government's Industrial Strategy Challenge Fund, Michael will assist over 110 researchers in creating new models for battery systems and researching next-generation, solid-state batteries.

UK Research and Innovation chief executive, Professor Sir Mark Walport, said, "This new supercomputer will be a valuable resource for the UK's battery researchers, providing them with the insight necessary to improve battery performance and lifetime, and reduce costs."

Developing more accurate simulations of batteries will give researchers the ability to design advanced batteries without the need to create prototypes for every new material.

The initial challenges to be tackled by the team include the fast-charging of batteries, low temperature operation and thermal management of cells within battery packs.

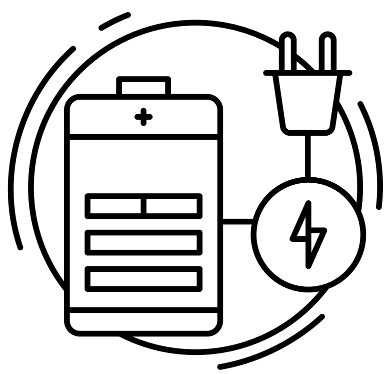

SUPERCOMPUTER TRAVELS TO SPACE

A novel 'space-qualified' supercomputer, developed by the University of Pittsburgh, set off on a journey to space on May 1, 2019.

This radiation-tolerant computer cluster, called the Spacecraft Supercomputing for Image and Video Processing (SSIVP) system, traveled 250 miles skyward with the SpaceX-17 mission to the International Space Station, where it will spend the next three to four years.

"[It will be] one of the most powerful space-qualified computers ever made and flown," said Professor Alan George, founder and director of the National Science Foundation Center for Space, High-performance, and Resilient Computing.

SSIVP, which is about the size of a breadbox, includes dual high-resolution cameras capable of snapping five-megapixel images of Earth.

Space is a challenging place for computers to thrive due to high fluctuations in temperatures, strong vibrations during launch and higher levels of radiation—all of which can affect performance.

During this time, the supercomputer will gather and monitor data on weather patterns, deforestation, the effects of natural disasters on Earth, and the effects of space and radiation on electronic devices.

Ultimately, the researchers' main objective is progression toward autonomous spacecraft, such as a more advanced version of the self-driving cars seen on Earth.

INDIA'S SUPERCOMPUTING MISSION INKS ATOS CONTRACT

India's Centre for Development of Advanced Computing (C-DAC) has entered into a three-year contract with French technology firm Atos, for the latter to supply its BullSequana supercomputers.

This contract is part of the National Supercomputing Mission (NSM), a seven-year plan of Rs 4,500 crore (~US$ 650 million) that aims to create a network of over 70 high-performance computing facilities across India (see *India's Need For Speed*, p. 38). Atos' BullSequana supercomputers have a cumulative computing power of more than 10 petaFLOPS and will be used by researchers in drug discovery, weather and natural disaster prediction, and even space exploration.

Although they will help India flex its computing muscle in various applications, the supercomputers are unlikely to be the fastest in the world.

"Since the beginning, [the computers built under] NSM were never targeted to be in the TOP500 list. We wanted to make the project based on applications that will be important for the country," Mr. Herman Darbari, director-general of C-DAC, told Indian business newspaper *BusinessLine*.

Digital Dispatch

SECOND ASIA-PACIFIC HPC-AI COMPETITION LAUNCHED

The second Asia Pacific High Performance Computing—Artificial Intelligence (HPC-AI) competition was launched on March 12, 2019.

Jointly organized by the HPC-AI Advisory Council and the National Supercomputing Centre (NSCC) Singapore, it encourages teams throughout the Asia-Pacific region to showcase their expertise in a friendly yet spirited competition that builds critical skills, professional relationships and lifelong comraderies.

In the 2018 inaugural edition, 18 teams hailing from seven countries, including China, Singapore, Japan, Korea and Thailand, participated in the competition. The winning team came from China's Tsinghua University.

This year, 20 teams have been tasked with achieving high-speed cosmological simulations of objects such as planets and galaxies with HPC resources. For the AI section of the competition, they are required to perform image recognition using a popular open-source AI software called TensorFlow.

The competition will continue through until July 26, 2019, and results will be announced in August. The winning team will be invited to participate in the 2020 International HPC-ISC Student Cluster Competition, to be held in Germany in June 2020.

SUPERCOMPUTING IN SINGAPORE GETS US$150M BOOST

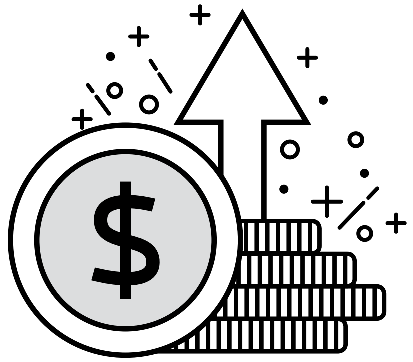

The National Research Foundation (NRF) of Singapore has set aside S$200 million (~US$146 million) to boost the nation's high performance computing capabilities.

Singapore already owns a one-petaFLOPS supercomputer—the Aspire 1—housed at the National Supercomputing Centre (NSCC) Singapore.

With the new funding, the capacity for high performance computation will increase by 15–20 petaFLOPS, supporting research on complex problems such as the modeling and prediction of climate patterns as well as the analysis of mobility conditions on Singapore's roads.

"Whether our national supercomputers are modeling climate change, or designing digital twins of manufactured products from tiny chips to the largest oil rigs, our research scientists and engineers are working together with our computer specialists to solve problems digitally and physically across industries," said Mr. Heng Swee Keat, then Singapore's minister for finance and NRF chairman.

COMPUTING THE M87 BLACK HOLE

An international team of researchers has revealed an image that many believed impossible to produce: a portrait of the shadow cast by a black hole that sits at the center of the galaxy Messier 87, which is 53.49 million light years away from Earth.

The image was captured using a distributed collection of eight high-altitude radio telescopes scattered around the globe, called the Event Horizon Telescope.

The simulations and visualizations, which were computationally expensive, required the support of three supercomputers at the Texas Advanced Computing Center (TACC)—Stampede1, Stampede2 and Jetstream.

Several teams of researchers used the high-performance computers to model the physical attributes of M87 and predict observational features of the black hole. These simulations ultimately gave scientists confidence that the image produced through observation matched up with reliable physical theories.

"These groundbreaking results show how discovery and imagination are fueled by advanced computing systems combined with the talented researchers and developers that use them," said Dr. Niall Gaffney, director of data intensive computing at TACC.

WHAT'S UP!

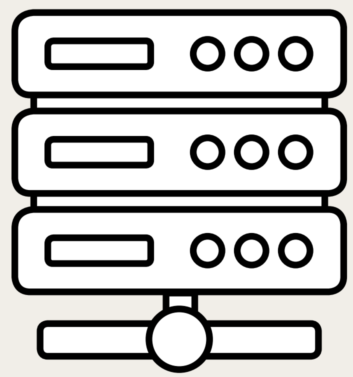

SC19 IS UP NEXT

The 2019 edition of The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC19) will be held at the Colorado Convention Centre in Denver, Colorado from November 17–22, 2019.

Since its inauguration in 1988, the SC Conference series' highly competitive technical program has never failed to showcase the latest innovations and breakthroughs pushing the frontiers in the world of HPC. In addition to programs such as talks, panels, tutorials, and research paper and poster presentations, SC19 features the "Birds of a Feather" session—a platform for attendees to openly discuss trending topics within the HPC community.

Another highlight is the SC Awards, which recognize outstanding researchers and their contributions to the field of HPC. This includes the prestigious ACM Gordon Bell Prize, and awards such as the IEEE-CS Seymour Cray Computer Engineering Award and the IEEE-CS Sidney Fernbach Memorial Award.

In addition, the Students@SC program specifically caters to undergraduate and postgraduate students. It offers mentorship matching, job fairs and HPC-related competitions to help students explore possible careers in HPC.

WHAT
SC19

WHEN
November 17–22, 2019

WHERE
Denver, Colorado, US

SUPERCOMPUTINGASIA 2020: SEE YOU THERE

After a successful 2019 edition that attracted over 700 delegates, Asia's premier HPC conference SupercomputingAsia (SCA) will be back next year from February 24–27, 2020, in Singapore. This year, SCA20 will focus on the theme 'HPC Powering Intelligent Cities.'

Organized by the National Supercomputing Centre (NSCC) Singapore, SCA20 will once again bring together the world's HPC leaders from both academia and industry. Delegates will gain access to insights shared by these thought leaders, as well as receive ample opportunities to network with the supercomputing community.

The annual conference will encompass an umbrella of notable supercomputing events, and delegates can expect an exciting calendar of keynote addresses, research presentations, workshops and industry-specific showcases. Save the date and stay tuned for more details!

WHAT
SupercomputingAsia 2020

WHEN
February 24–27, 2020

WHERE
Singapore

It's a MATERIAL WORLD

How high-performance computing is transforming the material world

Supercomputers are helping scientists run simulations, build databases and model chemical reactions, accelerating the discovery of new materials with novel applications.

By **Li Lidao**

Silicon—the backbone of modern electronics—was first discovered in 1824 by Swedish chemist Jöns Jacob Berzelius, largely via a trial and error process. Historically, this was how new materials came about: chemists in the lab would synthesize an entire collection of different substances before being able to find the one with all the properties they desired. It was, and often still is, an expensive and time-consuming process.

The advent of modern-day computing promises to change that. With the processing speed of the world's top supercomputers increasing by thousands of times over the past two decades, scientists can now perform incredibly complex calculations, based on fundamental equations of thermodynamics and quantum mechanics, to predict the behavior of both existing and new materials.

Supercomputing Asia showcases five materials designed by scientists using supercomputers and discusses how they could change the world.

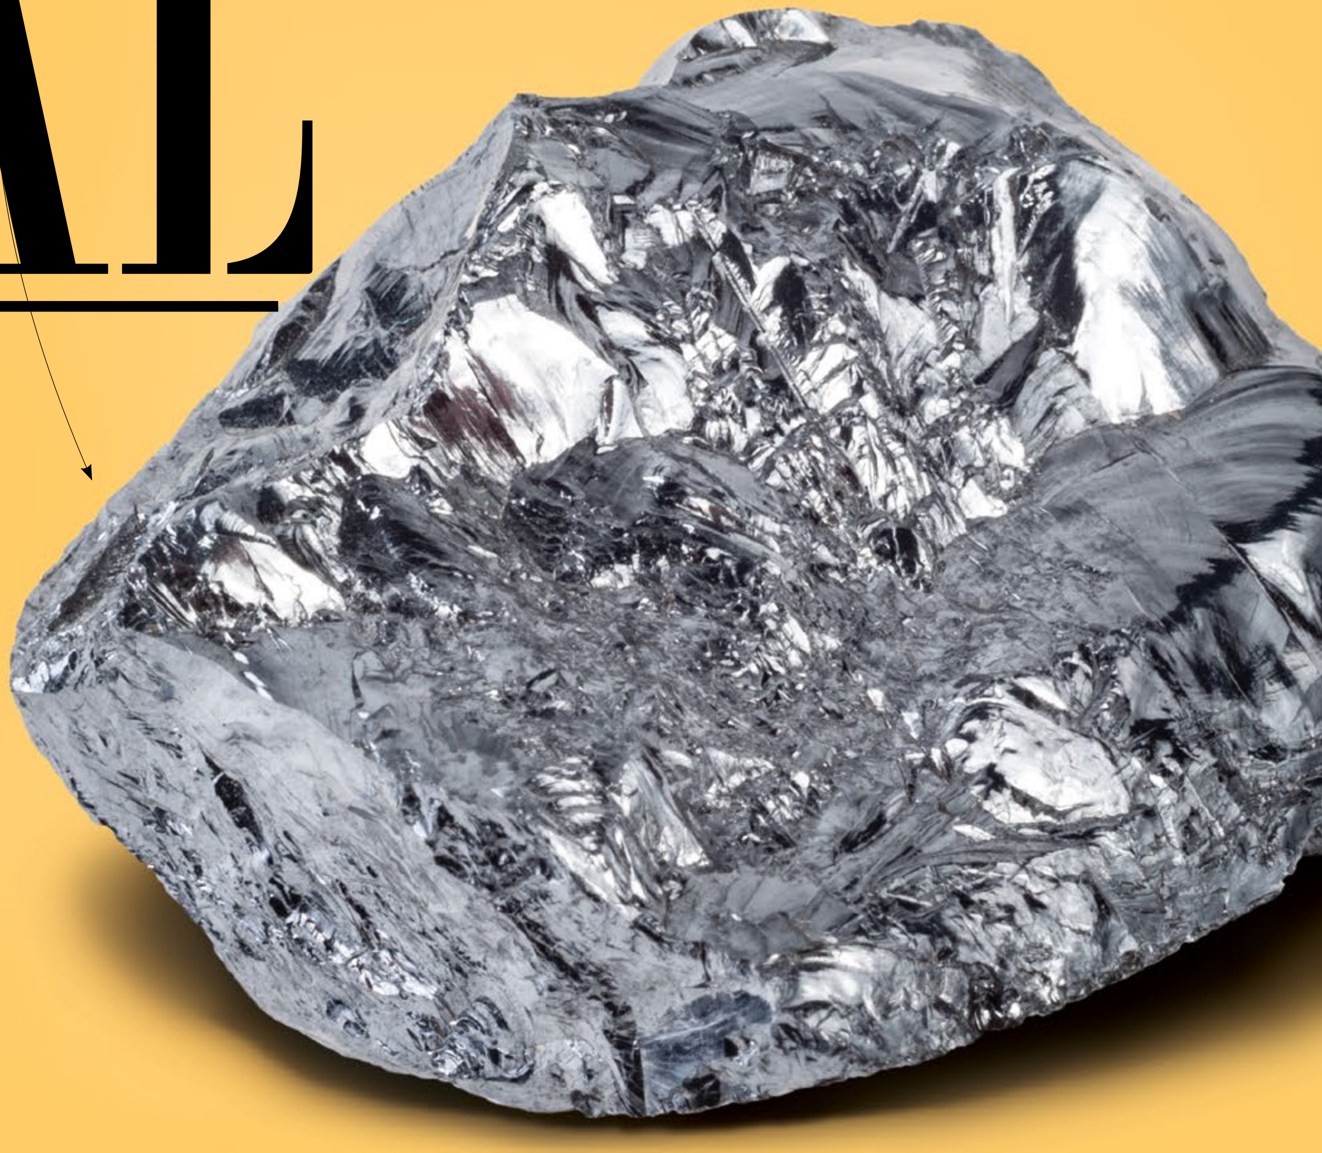

Alloys

Since ancient times, the creation of new alloys with superior properties has helped societies progress. When our predecessors mixed copper with tin, they formed an alloy that is harder than both original metals, marking the start of the Bronze Age.

Today, metallic glass is an emerging class of alloys well known for its great strength and malleability. Similar to conventional glass, it is easily molded to take on complex shapes when heated. Researchers are currently exploring how such materials can be shaped into biomedical implants that could replace bone tissue.

However, a myriad of alloys could theoretically be created from the 118 elements on the periodic table, and it is no easy feat to identify metallic glass from that list. This is where supercomputers could help scientists identify the most likely combinations as candidates for metallic glass.

Using a Hitachi SR11000 supercomputer at Tohoku University in Japan, a team of researchers simulated the arrangement of atoms in a well-known example of metallic glass. With a peak performance at over seven teraFLOPS, the supercomputer's computational power was crucial to determining how the alloy's atomic structure gave it its glass-forming ability. This theoretical understanding could aid in the design of novel alloys with similar properties.

Scientists are using supercomputers to investigate the properties of metallic glass, durable and moldable alloys that could be used as bone implants.

Powdered radioactive fuel is oxidized into ceramic pellets like these. Ceramics are also used in nuclear reactor cores and to store nuclear waste.

Ceramics

Forged at temperatures exceeding 2,500 °C, ceramics are well known for their chemical inertness and tolerance of harsh conditions such as high temperature and pressure. While they are often associated with teapots and tableware, ceramics have found a niche in nuclear power generation, as a component of the nuclear reactor core, or as a matrix for nuclear waste storage.

Although it is critical to enhance the radiation tolerance of ceramic materials to increase their life span and reliability, such experiments are challenging to perform due to the toxicity of the radioactive materials involved. By simulating the effects of radiation on ceramic structure, supercomputing thus emerges as an attractive and safe alternative to conducting experiments.

One such experiment was performed on the Tianhe-1 supercomputer located at the National Supercomputer Center in Tianjin, China. Leveraging its 1.2 petaFLOPS peak computing speed, the research team simulated how radiation could knock off atoms in ceramic structures to cause defects, finding that prior structural misalignments known as stacking faults actually reduce the damage to ceramics from radiation. This knowledge paves the way for the design of ceramic materials that are more resilient under radioactive exposure.

Polymers

Waterproof and slow to degrade, plastics have earned a bad reputation for contaminating ecosystems and endangering biodiversity. This much maligned material belongs to a class of compounds known as polymers, which vary in physical properties such as stiffness, strength and melting point.

The manner in which the individual repeating units of a polymer interact with one another gives the resultant product its unique properties. Scientists therefore yearn for greater control over these interactions so that they can precisely engineer polymers for various applications.

At the Indian Institute of Technology-Guwahati in India, scientists are using the Param-Ishan supercomputer to simulate the best way to synthesize a super hydrophobic—or 'water-fearing'—polymer that could be used to clean up oil spills. They screened five different ions as potential catalysts for the polymerization reaction and found lithium to be most effective due to its small atomic size. This research resulted in the development of a polymer that could soak up ten times its weight of motor oil.

Another widely-used polymer is rubber, and Japanese company Sumitomo Rubber Industries Ltd is using the K supercomputer at the RIKEN Advanced Institute for Computational Science in Japan to simulate how rubber and silica particles in automobile tires interact when mechanical stress is applied as part of their daily operation. This information allowed the company to develop safer, more durable and more fuel-efficient tires.

While plastics can be harmful to marine life and humans, other types of polymers can be used to clean up oil spills and make more fuel-efficient tires.

Semiconductors

The Samsung Galaxy Fold, a smartphone that unfolds into a tablet, has given us a glimpse into a future where 'bendy' electronic devices are possible. But how does the screen of such a phone bend without breaking?

Organic light emitting diodes (OLEDs), which are made from organic semiconductors and give off light when an electrical current is applied, are a key material that makes such flexible displays possible.

For OLEDs to be stable and efficient, the molecules inside them must adopt an ordered arrangement. To optimize the conditions for achieving this ordered arrangement, Samsung relied on simulations performed on the Tachyon II supercomputer at the Korea Institute of Science and Technology Information's supercomputing center in South Korea. The scientists found that lower temperatures and the use of longer, more planar organic semiconductor molecules were ideal for creating flexible OLED displays.

Semiconductors are also important in the fabrication of solar cells used in photovoltaic sytsems. Once again, the structures and properties of the organic semiconducting layer on the surface of solar cells are being modeled using supercomputers such as the IBM Summit at Oak Ridge National Laboratory, the fastest supercomputer in the world at the time of writing, with a peak performance of 200 petaFLOPS.

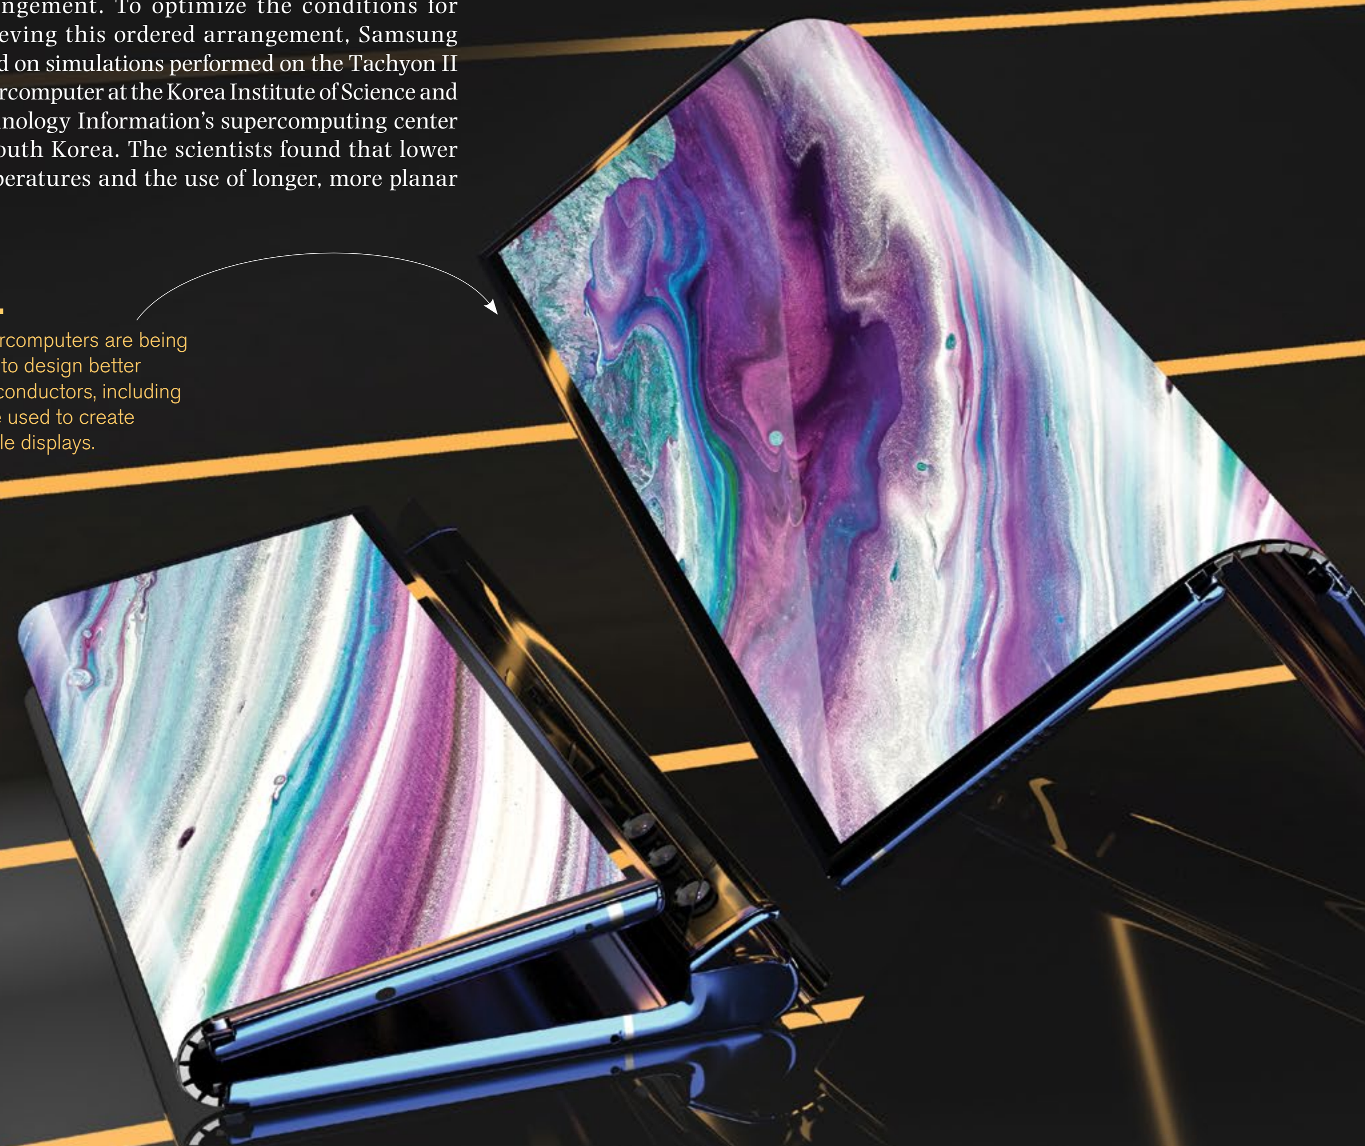

Supercomputers are being used to design better semiconductors, including those used to create flexible displays.

Quantum materials

To make quantum computers a reality, new materials that allow scientists to control electrons on a quantum scale are urgently needed.

IBM Q, the world's first commercial quantum computer, has its processors kept at an unworldly temperature of close to absolute zero. This is necessary to prevent external interference from disrupting the quantum states of the particles—known as quantum bits, or qubits—used to store information. Even so, the duration of information retention does not last beyond the range of microseconds, and the high error rate this causes remains one of the main challenges facing quantum computing today.

To overcome the limitations of existing qubits, scientist are searching for novel materials to create 'spin qubits' which promise to reduce the error rate of future quantum computers. For instance, researchers used the Magnus Cray XC40 supercomputer at the Pawsey Supercomputing Centre in Australia, which has a peak performance exceeding one petaFLOPS, to study how electrons interact with each other in a silicon donor-quantum dot system.

The computational results helped to illustrate how the voltage could be tuned to compensate for potential errors in their proposed quantum computing system. This knowledge paves the way for building more stable and functional quantum computers that may one day also have a role to play in chemistry research and materials design.

1,000,000,000,000,000,000

THE RACE TO EXASCALE

GETTING ARTIFICIAL INTELLIGENCE AND BIG DATA UP TO SPEED

Professor Lu Yutong shares how China's past experiences have shaped the country's supercomputing efforts, and calls for more collaboration among Asian countries.

By **Rebecca Tan**

The Tianhe-1A supercomputer made history when it was unveiled in 2010 as the first Chinese supercomputer to headline the TOP500 list of the world's fastest supercomputers. For the first time since the biannual list was published in 1993, a Chinese supercomputer was recognized as the most powerful in the world, displacing long-time leaders like the US and Japan.

In the decade since, continued Chinese dominance of the TOP500 list has proven that the feat was no flash in the pan but a milestone in China's high performance computing (HPC) journey. Tianhe-1A's achievement was repeated in 2013 by its successor—Tianhe-2—and once again in 2016 by Sunway TaihuLight, enabling China to claim bragging rights as owner of the world's most powerful supercomputer for five years in a row (see *A Decade of Progress*, p. 22).

While China's rise might appear unusually rapid—the country did not have a system powerful enough to make it to the top ten till 2010 and only began appearing in the TOP500 at all in 2001—its current supercomputing prowess is actually the result of decades of work behind the scenes.

To understand how Chinese supercomputing has developed and where it is headed in the future, we spoke to Professor Lu Yutong, director of the National Supercomputing Center in Guangzhou, China, at the SupercomputingAsia conference held in Singapore from March 11-14, 2019.

FROM GENERATION TO GENERATION

One of the few people to have extensive first-hand knowledge of the successive generations of Chinese supercomputers, Lu traces her academic lineage to National University of Defense Technology (NUDT), home to China's first supercomputer—Yinhe-1—which was launched in 1983. As a computer science undergraduate, Lu worked on compiler software for Yinhe-2, developing a taste for writing code and going on to complete both her masters and doctorate degrees at the same institution. "In a way, I've grown up with the Chinese domestic supercomputer," Lu shared.

At NUDT, Lu focused on designing and implementing supercomputers, rising through the ranks to become the design director of the Tianhe-1A and Tianhe-2 systems, both of which held number-one positions on the TOP500 list. Apart from cementing China's place in the international supercomputing scene, the Tianhe series also represented a significant shift in supercomputer architecture trends—from using ever more general purpose central processing units (CPUs) to a mix of CPUs and accelerators.

"The Tianhe-1 was the first heterogeneous large-scale supercomputer in the world, using a mix of CPUs and graphics processing units (GPUs)," Lu explained. "At that time, many scholars asked me why we chose the heterogeneous approach as it was relatively hard to use. We did indeed have a difficult time optimizing the applications on Tianhe-1 as the GPU part was not as mature as it is today; we had to work more on the software, libraries and data transfer optimization."

"Supercomputer design, however, is a trade-off," she continued, explaining how using a heterogeneous architecture allowed the Tianhe team to achieve performance while balancing equally important considerations such as power consumption, system cost and footprint. Trends in the TOP500 have since validated their approach, Lu said, with heterogeneous systems now making up the majority of the most powerful systems on the list in the last ten years.

DOMESTIC DEVELOPMENTS

China's remarkable progress on the Tianhe systems did not go unnoticed by the international community, and not all reactions were positive. Alarmed by the country's newfound potential to use supercomputers for the design of nuclear weapons, in

> "EVERY COUNTRY WANTS TO GET TO EXASCALE NOT JUST TO BE THE FIRST BUT BECAUSE OF THE REQUIREMENTS OF SCIENTIFIC RESEARCH, TECHNOLOGICAL INNOVATION AND INDUSTRY.
>
> EXASCALE IS NOT THE END GOAL BUT A MEANS TO AN END."

PROFESSOR LU YUTONG

ALL SIGNS SUGGEST THAT KEY APPLICATIONS FOR FUTURE EXASCALE SYSTEMS ARE BIG DATA AND ARTIFICIAL INTELLIGENCE.

2015 the US banned the sale of Intel chips to China. However, rather than stymying China's advance, the move is largely considered to have backfired, accelerating the development of domestically-produced chips instead.

Just one year after the Intel ban, China's Sunway TaihuLight stormed to the first place on the TOP500 list with a peak performance 125 petaFLOPS, more than double the previous record holder Tianhe-2's 54 petaFLOPS. But more than its impressive performance, Sunway TaihuLight stood out for being a fully made-in-China or 'self-controllable' machine, featuring not only custom interconnects and memory but also the Shenwei SW26010 processor that was designed and fabricated domestically. The message was clear: China was determined to succeed at supercomputing, with or without the help of US chips.

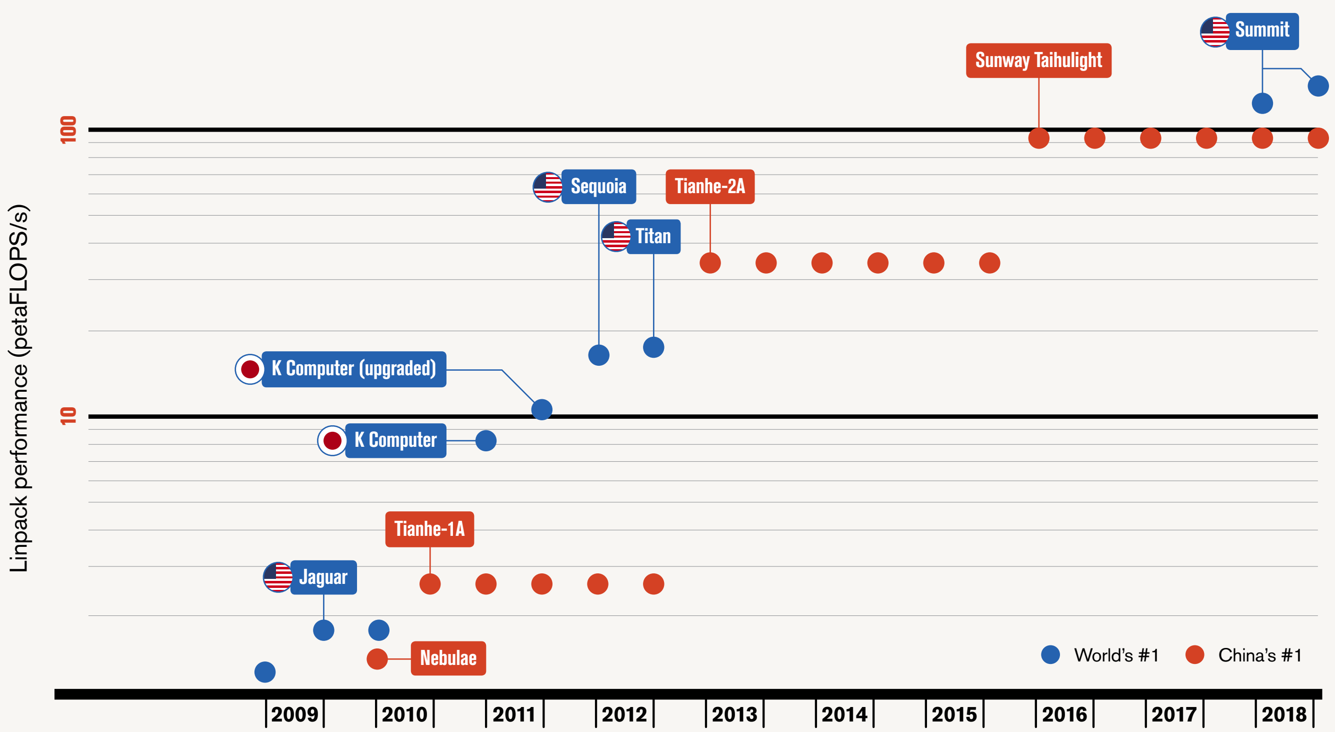

A DECADE OF PROGRESS

In 2009, China was not considered a major supercomputing superpower, and had no entries in the top ten of the TOP500 list. However, by June 2010 the Nebulae supercomputer at the National Supercomputing Center in Shenzhen was ranked second in the world, and in November 2010 the Tianhe-1A made its historic appearance at the top of the list. Although Tianhe-1A was quickly eclipsed by Japan's K computer, China went on to enjoy a five-year winning streak with Tianhe-2 and Sunway TaihuLight, only ceding to the US' Summit supercomputer in 2018.

That year, three out of six finalists for the Gordon Bell Prize—which honors the most outstanding HPC applications—were from China, and the highly coveted award was ultimately given to a team that ran a weather simulation over ten million cores of Sunway TaihuLight, definitively demonstrating that it was no mere 'stunt machine' but a fully functional, even prize-winning, one.

This application-centric focus has also guided Lu in her role as director of the National Supercomputing Center in Guangzhou. For her, while designing and deploying increasingly powerful systems remain important, developing a flexible and easy-to-use interface is a top priority. "Many of our users are not computer science experts but domain-specific scientists so they may not be very familiar with HPC programming models such as message passing interface (MPI)," Lu said. "We have tried to design a platform that can be used by scientists and even industry users, providing some domain-specific custom applications to help them."

CHINA'S THREE-HORSE RACE TO EXASCALE

China's reign at the peak of the TOP500 has recently been disrupted by the US, whose Summit and Sierra systems now hold first and second position respectively on the latest edition of the list. Together with the Perlmutter supercomputer, which is expected to be deployed at the Lawrence Livermore National Laboratory in 2020, Summit and Sierra are stepping stones towards the goal of exascale supercomputers, machines that can perform a billion billion (10^{18}) calculations per second. To date, the US has set aside US$430 million for their Exascale Computing Project, and at least US$400 million for the Aurora exascale supercomputer, expected to come online in 2021.

China, too, has made considerable investments into developing exascale or e-class supercomputers, and is a serious contender in the race to exascale. Three teams—namely Sunway, Sugon and Tianhe—have successfully deployed pre-exascale prototypes, each with its own approach drawn from academia, industry and running a national supercomputing center, respectively. While the teams may vary in their architectural approaches, all three prototypes make use of self-controlled chips, underscoring the importance of self-reliance to the Chinese.

Lu leads the Tianhe team, which completed their prototype in July 2018. Installed at the National Supercomputing Center in Tianjin, the Tianhe prototype is a heterogeneous system that relies on an undisclosed processor paired with an improved version of the Matrix-2000 accelerator used in Tianhe-2A. The Sugon prototype is similarly heterogeneous, making use of a Hygon processor and accelerator, while the Sunway prototype features the SW26010 chips used in TaihuLight.

"Even in the US, Japan and Europe, people are still not sure which architecture will suit their priorities best, so we need different approaches," Lu said. "For Tianhe, we will have more users, not only numerically but also in terms of the range of application areas. Our goal is to establish a general purpose exascale system rather than a specialized one, so we have a different technical approach from the other two teams."

The strategy of investing in different architectures in parallel is designed to deepen and broaden Chinese expertise, with knowledge gleaned from all three approaches expected to contribute to the final exascale system to be deployed between 2020 and 2021. "Nothing will be wasted," Lu added.

THE FUTURE IS INTELLIGENT

Regardless of which approach will be implemented internally in China or by the US, Japan or Europe, all signs suggest that key applications for future exascale systems are big data and artificial intelligence (AI). Supercomputers have traditionally been used for HPC applications such as climate modeling and scientific computation, and while these workloads will likely remain a mainstay of supercomputing centers around the world, there is increasing convergence between HPC and big data/AI applications.

Last year, Japan launched the AI Bridging Cloud Infrastructure (ABCI), a supercomputer purpose-built for AI and machine learning applications on the cloud. Similarly, the Aurora exascale supercomputer has been earmarked for AI projects in neuroscience and personalized medicine, among others. China, meanwhile, has identified deep learning applications such as tumor diagnosis and video analytics as key focus areas.

"Every country wants to get to exascale not just to be the first but because of the requirements of scientific research, technological innovation and industry; the world needs computing power. In other words, exascale is not the end goal but a means to an end," Lu said.

"Over the past 40 years we have seen very clearly how supercomputing leads technology development for the whole IT industry," she continued, citing the example of how China's decades-long investment in supercomputing has paid innovation dividends for its domestic IT sector with success stories like Sugon, Inspur and Huawei.

Investment in supercomputing capabilities is particularly strategic for Asia, Lu added, noting that Asia has a high demand for HPC simply by virtue of its large population. "I hope to see more collaboration between Asian countries, particularly more economically developed ones like Japan, Korea and Singapore; we need to work together to boost the HPC community in Asia."

1,000,00
000,000

A QUANTUM LEAP INTO THE FUTURE

Building an 'unhackable' internet

By harnessing the unique properties of quantum mechanics, scientists are developing radically new technologies to make the internet impenetrable to hackers.

By **Yvonne Gao**

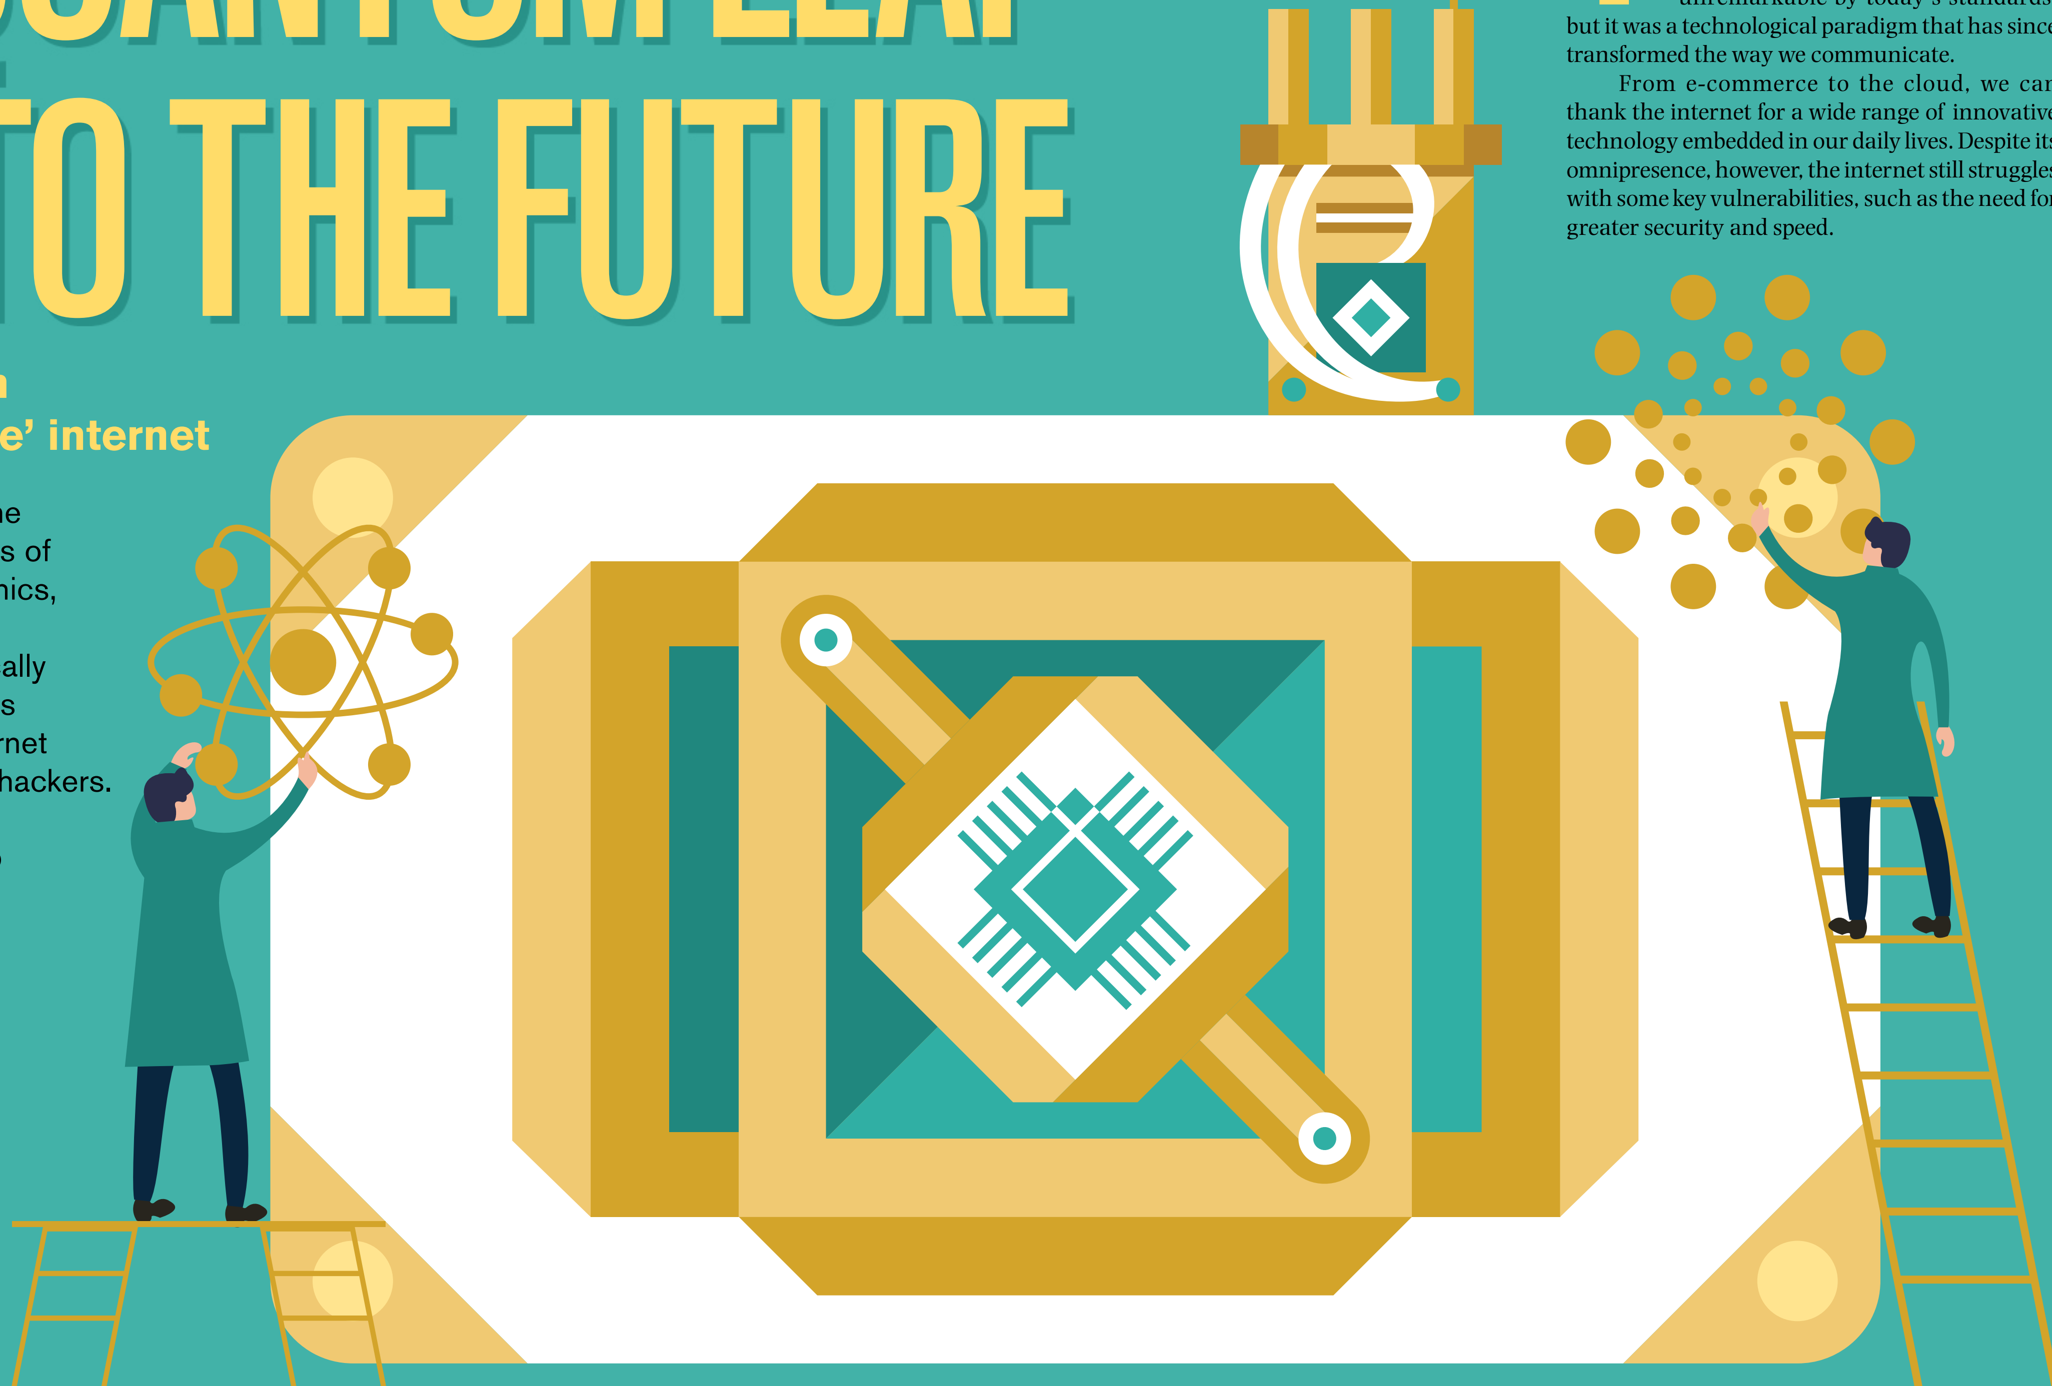

This year marks the 50th anniversary of the first demonstration of information transfer between two computers on ARPAnet, the grandfather of today's internet. The simple one-word message—LOGIN—communicated on the fledgling ARPAnet may seem rather unremarkable by today's standards, but it was a technological paradigm that has since transformed the way we communicate.

From e-commerce to the cloud, we can thank the internet for a wide range of innovative technology embedded in our daily lives. Despite its omnipresence, however, the internet still struggles with some key vulnerabilities, such as the need for greater security and speed.

To address the internet's Achilles' heel, scientists have turned to quantum physics for its unique and peculiar properties—quantum bits of information, or qubits, that can simultaneously live in a quantum superposition of 0 and 1, in contrast to classical bits of information that are strictly binary.

A quantum computer therefore can solve complex problems much faster than a classical computer, and a quantum communication network has the potential to bring unprecedented capacities for transmitting and processing large volumes of data. Not only does it promise improved speed, a quantum internet can also transmit large amounts of encrypted data, making it unhackable.

COMMUNICATION IN A FLASH

In a quantum network, the qubits are typically constructed using quantum states of photons, or single particles of light, said Assistant Professor Gao Weibo, an experimental physicist studying photonic quantum computation and communication at Nanyang Technological University, Singapore. Gao was speaking at the recent SupercomputingAsia 2019 conference track on the latest in quantum communication research.

"Photons are great carriers of quantum information because we can control them very precisely. They are also ideal for building a quantum communication network, as they can be coherently transmitted over very long distances," Gao said.

Gao is part of a larger team led by Professor Pan Jianwei of the University of Science and Technology of China, who is frequently referred to as "the father of quantum." Under Pan's leadership, the researchers are feverishly developing technologies that will someday form the backbone of a global quantum network.

"For us to build a quantum network, it will require both the interference and entanglement of individual photons," said Gao. "Hence, it is important for us to build high-quality single photon sources and efficient photon detectors."

SPOOKY CYBERSECURITY

Compared to the traditional communication networks that are prone to malicious eavesdroppers, a quantum network exploits the unique properties of quantum physics to provide an unbreakable system of information transfer.

The key element that offers protection against interception is something called entanglement, undoubtedly one of the most iconic and intriguing concepts in quantum mechanics. Albert Einstein himself once described entanglement as "spooky action at a distance," for good reason.

"When two photons are entangled, their properties become so inextricably intertwined that any minute external influences experienced by one of them is instantaneously 'felt' by its twin, even if they are placed at two opposite ends of the universe," Gao explained.

As bizarre as it sounds, entangled photon pairs are critical ingredients for making a secure quantum network possible. For instance, by distributing half of an entangled photon pair to each node of the network, two parties can exchange confidential information by simply implementing specific operations on their half of the entangled pair.

"If there is an eavesdropper, the quality of the entanglement correlations will be irreversibly degraded, allowing us to conclusively certify whether there has been an attempt to 'steal' the information," said Associate Professor Alex Ling, a principal investigator at the Centre for Quantum Technologies in Singapore, also speaking at SupercomputingAsia 2019.

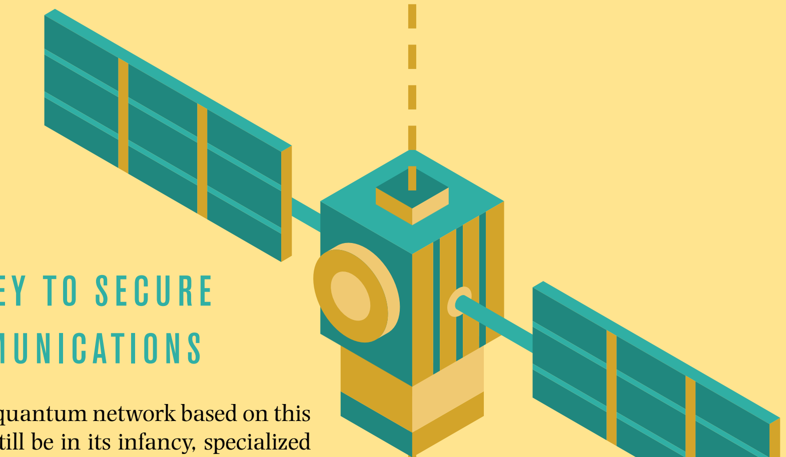

THE KEY TO SECURE COMMUNICATIONS

While a full-blown quantum network based on this technology might still be in its infancy, specialized functionalities are already within reach today. One such near-term application is quantum key distribution (QKD), which uses entangled photons to provide a shared cryptographic key for two parties to securely exchange messages over conventional networks.

What QKD ensures is that the quantum network will remain robust against future technologies, even the omnipotent quantum computer which threatens to break many of the conventional encryption schemes in a matter of minutes.

"Although the concept of creating entangled photons seems rather complicated, the technology is actually quite mature and accessible now," Ling said. "We have started to deploy such entangled photon sources in many different environments outside the laboratory. In fact, one of our systems is currently housed in a shoebox-sized satellite to be launched in collaboration with the Japanese space agency, JAXA."

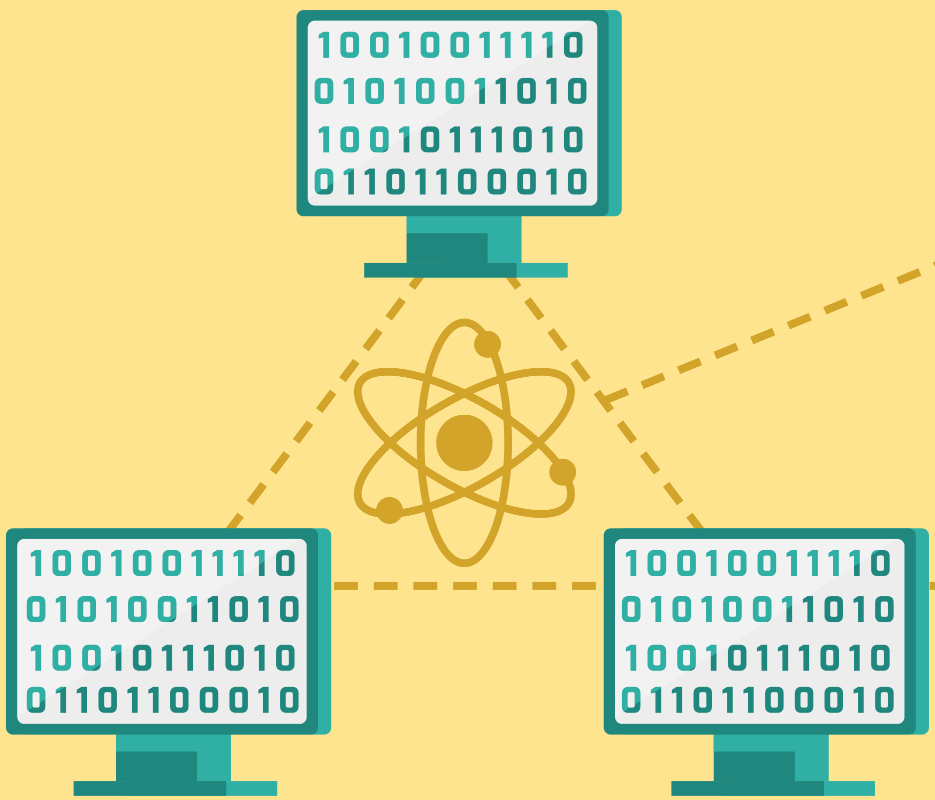

Thanks to experimental advancements over a decade, standalone QKD devices can now be easily purchased off the shelf. Quantum networks capable of implementing QKD protocols have also been steadily expanding, both in distance covered and generation rates.

"Progress has been further accelerated by the development of satellite-based QKD schemes, because photons travel over much longer distances through space without getting lost," Gao added.

Indeed, China's Micius satellite enabled in 2017 one of the most spectacular demonstrations of QKD technology, by establishing the very first intercontinental video call between Beijing and Vienna that was thoroughly unhackable.

RE-IMAGINING THE INTERNET

As we continue to push the boundaries of quantum technologies, the ambition of a large-scale quantum internet is turning into reality. Near-term applications have already been identified, such as precise clock synchronization to efficient verifications of distributed data.

"We can also realize high-precision metrology by exploiting entanglement correlations distributed over large distances. That is why it is so exciting to work with these technologies—there are so many exciting applications to look forward to," said Ling. "As long as we are sufficiently careful with the human artifacts in the system, the security of such a quantum network can be inherently guaranteed against computational advances."

Looking further down the road, one can imagine a robust quantum Internet of Things where quantum processors and sensors are securely connected to one another, achieving unparalleled capabilities that are impossible with classical systems.

Although we may not be able to predict exactly when a mature quantum internet would be available for all, one thing remains certain—it will completely transform the way we think about communication, just as the internet did decades ago.

HOW DOES QUANTUM KEY DISTRIBUTION WORK?

Quantum key distribution, or QKD, uses entangled photons to provide a shared cryptographic key for two parties to securely exchange messages over conventional networks.

1 A pair of entangled photons is created by driving a non-linear crystal with a very powerful laser. An entangled pair refers to two particles of light whose properties are so inextricably linked such that changes made to one can be instantaneously experienced by its twin.

2 Each half of the entangled pair is distributed to one end of the communication channel. This can be accomplished by using highly focused beams from a satellite.

3 Alice and Bob, the two parties who wish to communicate secretly, now each have half of the entangled pair. They can compare the quantum states of their respective halves via simple measurements to verify if an interception has been attempted.

4 After the verification, Alice and Bob use the entangled photons to create a secret key and encrypt their data with it.

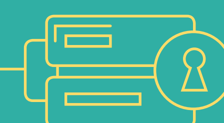

5 Alice and Bob are now able to communicate encrypted data securely via conventional networks without having to worry about potential eavesdroppers.

MOVE IT!

Helping Big Data go round the globe

Teams raced to move two terabytes of data across five different countries in the inaugural 'Move that Data!' Data Mover Challenge.

By **Li Lidao**

As compared to a time when goods were moved between countries on horse carts, the volume of cargo that can cross geographical boundaries has increased exponentially with the invention of container ships and cargo planes. Analogous to physical cargo, Big Data must also be moved around the world quickly. Therefore, innovations for borderless, rapid data transfer are urgently needed.

To encourage data scientists and researchers to think up novel data transfer strategies, the National Supercomputing Centre (NSCC) Singapore organized the inaugural 'Move that Data!' Data Mover Challenge.

Seven teams participated and were each given one week to deploy software to transfer two terabytes of data across five different countries, namely NSCC in Singapore, National Computational Infrastructure (NCI) in Australia, Korea Institute of Science and Technology Information (KISTI) in Korea, National Institute of Information and Communications Technology (NICT) in Japan, and StarLight in the US.

Supercomputing Asia caught up with one of the organizers, Associate Professor Francis Lee Bu Sung of Nanyang Technological University, Singapore, and vice-president of the Singapore Advanced Research and Education Network (SingAREN), to find out more about the challenge.

Why is high speed data transfer important for research?

Assoc. Prof. Francis Lee: In this day and age of digital data and data-intensive computing, moving scientific data from one spot to another is important for global collaboration. If you cannot move data fast enough, scientists working together on a research project will not be able to make their discoveries in a timely manner.

In the past, data was moved around the globe by FedEx, which means there was a lag between data generation and data analysis. But things have changed; the volume of data has grown even bigger, and people want things faster online, even on the spot. I have seen some collaborations among researchers where the minute they capture data, it is shared with other researchers elsewhere in the world.

We are tackling bigger problems with bigger data, so I think rapid data transfer is very essential. There is no use having to wait one day or two for data to be transferred; you need to get it out as fast as possible.

Why is transferring research data more difficult than transferring everyday internet data?

FL: The transfer of everyday internet data, like streaming data and such, usually involves what we call small files, within the size range of a few hundred of megabytes to gigabytes. On the other hand, research data—like genomic data, for instance—is as large as one terabyte!

Even when data files are not that large, transferring large volumes of data at a time can be a challenge. For example, I might be helping Indonesia with the transfer of satellite data which may not be as large as one terabyte, but which consists of a lot of files, each of them a satellite image between 60 and 200 megabytes in size.

To enable the fast transfer of such research data, we need to tune the device parameters of data transfer nodes, such as the network card parameters, the input/output (I/O) scheduler, disk access and so on. Just by tuning the I/O scheduler, you can optimize the sequence in which data is read into the computer and sent out to the network. This allows you to more than double the throughput and achieve higher speeds of data transfer needed to move large amounts of data.

Francis Lee Bu Sung
Nanyang Technological University, Singapore
Singapore Advanced Research and Education Network

WINNERS OF THE 'MOVE THAT DATA!' DATA MOVER CHALLENGE 2019:

Overall Winner: Zettar Inc.

Project title: Zettar zx hyperscale data distribution software platform

Using their software, the team managed to achieve an amazing 68 Gbps transfer rate between Chicago and Singapore

Most Innovative Award: StarLight/iCAIR (International Center for Advanced Internet Research) team

Project title: STARLIGHT DTN-as-a-Service for Intensive Science

The team impressed the judges with a flexible framework involving various modules to supports large data transfers

What are the motivations behind organizing this challenge?

FL: Singapore has a data transfer rate of 100 Gbps to Japan and the US, and very soon to Europe as well. But it is not just about the bandwidth, you also need software that can make use of that bandwidth. That's why the Data Mover Challenge came about. We wanted to tackle the question: How can we get the best tools to move data efficiently from one point to another, so they can be shared?

At SingAREN, we have been tuning our servers and doing some data transfer optimization work as well, but we felt that if we opened up and engaged with external stakeholders in the field, we could learn and adopt best practices that we may not have thought about. And true enough, we found people doing things in innovative ways.

I think the main advantage of throwing out a challenge like this is the diversity of participants we get to engage with. Academics and industry players alike came in to try and solve something that is relevant to everyone in the field. We had teams from Argonne National Laboratory and FermiLab in the US who took part alongside others from the University of Tokyo and the Japan Aerospace Exploration Agency.

From the private sector, companies like Fujitsu and Zettar joined in as well, so we really had a good range of expertise involved. More importantly, because everyone came under the same roof at the SupercomputingAsia 2019 conference, there were many opportunities to interact and exchange notes, and hopefully take the field of rapid data transfer to the next level.

What are some of the hurdles the teams had to overcome?

FL: Each team was given one week to set up their software in multiple nodes in Japan, Korea, the US, Singapore and Australia, then demonstrate the transfer of two terabytes of data disk-to-disk, not just memory-to-memory. Eventually, our intention is to deploy good data transfer node software around the globe, which is why deployment across the five sites is essential.

There was a variety of large and small files to be transferred, thus allowing us to assess how the teams' software handled different packets of data. We gave participants a week to set up, after which the data transfers were carried out while we monitored the process. Thereafter, the teams had to present to the judges their data transfer rate and how they achieve that rate.

While the speed of data transfer was a main criterion for judging, it was not the only one. For example, we also wanted to know how each team's software maintained the data transfer rate over international links when the delay among those links is different. You need to demonstrate what the performance is when the delay is very large, when the delay is not so large, and so on. It takes a lot of effort to deploy software over five servers, make them run properly and tune them accordingly.

What skills did the teams need to have in order to succeed?

FL: I would say that a variety of skills are required, from computer science to network engineering. Some participants had software development backgrounds while others were good at tuning the networks. Of course, teamwork was essential as well.

Overall, I think all the teams performed very well. Initially, we said we only wanted to announce one overall winner. But as the competition went along, we felt that there was another deserving team, so we decided to recognize that team—the StarLight/International Center for Advanced Internet Research team—with the 'Most Innovative' award.

At the end of the day, we hope that competitions such as this will facilitate the transfer of ideas, knowledge and skills among people from diverse disciplines to enable faster transfer of research data around the world.

SCALING NEW HEIGHTS AT SCA19

Highlights from SupercomputingAsia 2019

SupercomputingAsia 2019 (SCA19) saw more than 700 research and industry delegates from around the world attending the three-day conference, which took place from March 12 to 14, 2019, in Singapore.

Graced by guest of honour Mr. Heng Swee Keat, deputy prime minister of Singapore, the conference ran under the theme of "HPC Futures – Hyperscalers, Exa, AI, Quantum and Beyond."

Program highlights included sessions on quantum computing, precision medicine and hyperscale infrastructure. Co-located events with SCA19 included the Conference on Next Generation Arithmetic (CoNGA), the Asia Pacific Research Platform (APRP) Conference, the Singapore-Japan Joint HPC Session, Supercomputing Frontiers Asia (SCFA), and the ASEAN HPC Workshop.

Winners of this year's SCA Awards were Mr. Lim Soon Hock, founder and managing director of PLAN-B ICAG, who won the SCA19–Singapore HPC-Pioneer & Achievement Award, while Professor Satoshi Matsuoka, director of the RIKEN Center for Computational Science, won the SCA19–Asia HPC Leadership Award. Finally, Dr. Raj Thampuran, special advisor to the Agency for Science, Technology and Research, won the SCA19–Asia HPC Visionary Award.


2


3


5

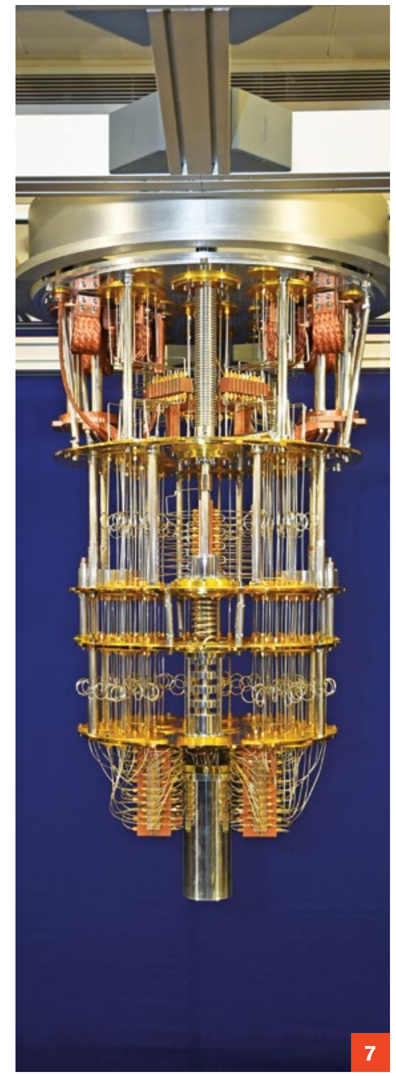
7

1


4


6

1. It was an occasion to celebrate for the three winners of the SCA19 awards, who were nominated and selected by a committee of distinguished international HPC experts.
2. Guest of honor Mr. Heng Swee Keat, deputy prime minister of Singapore, gave the opening address at SCA19.
3. Dr. Christine Ouyang, distinguished engineer at IBM Quantum Computing Technical Partnership and Systems Strategy, IBM Master Inventor, and member of Academy of Technology, giving the opening keynote, "From sci-fi to reality: quantum computing is here."
4. Professor Lin Dahua, co-founder of SenseTime and assistant professor at The Chinese University of Hong Kong, giving a plenary keynote.
5. From left to right: Mr. Peter Ho, chairman of the National Supercomputing Centre (NSCC) Singapore steering committee; Professor Satoshi Matsuoka, director of the RIKEN Center for Computational Science; Deputy Prime Minister Heng Swee Keat; Mr. Lim Soon Hock, founder and managing director of PLAN-B ICAG; and Dr. Raj Thampuran, special advisor to the Agency for Science, Technology and Research.
6. Deputy Prime Minister Heng Swee Keat taking a we-fie with SCA19 speakers and delegates in front of the IBM Q display.
7. A model of the IBM Q quantum computer.

8. Prof. Satoshi Matsuoka giving an update on supercomputing activities taking place in Japan.

9. From left to right: Mr. Andrew Howard, cloud team manager of National Computational Infrastructure Australia; Dr. Chin Fang, founder & CEO of Zettar Inc.; and and Associate Professor Tan Tin Wee, chief executive of NSCC, presenting an award to Zettar Inc. for being the overall winner of the Data Mover Challenge 2019.

10. Delegates from Fujitsu.

11. The SCA19 organizers giving a toast at the networking dinner and award ceremony.


8

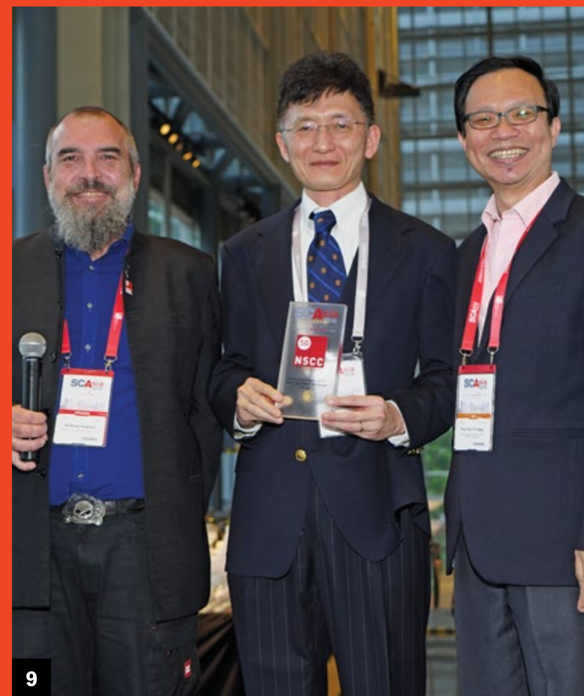
9


10


11

12


13

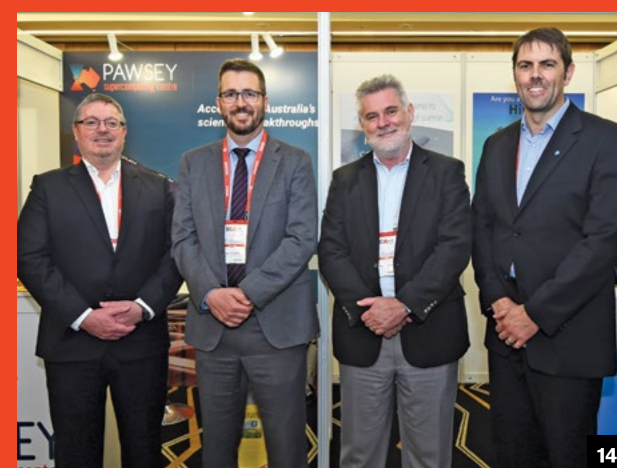

14


15

16

12. The organizers taking a group shot after another successful run of SupercomputingAsia.

13. Delegates from Alibaba Cloud.

14. Delegates from Pawsey Supercomputing Centre and the Commonwealth Scientific and Industrial Research Organisation (CSIRO) with the Australian High Commissioner to Singapore, His Excellency Mr. Bruce Gosper (far left).

15. Professor Lu Yutong, director of the National Supercomputer Center in Guangzhou, giving a plenary keynote.

16. See you at SupercomputingAsia 2020!

Business Bytes

FUJITSU TO BUILD JAPAN'S FIRST EXASCALE COMPUTER, POST-K

Fujitsu announced that it has completed the design of Japan's first exascale supercomputer, Post-K. Performing computations at one billion billion FLOPS per second, Post-K is expected to run one hundred times faster than its predecessor, K supercomputer.

Post-K will be equipped with A64FX™, a high-performance CPU developed by Fujitsu using Arm instruction set architecture. After detailed designs and system prototypes, the project cleared an interim evaluation by Japan's Council for Science, Technology, and Innovation.

In the next stage, Fujitsu and RIKEN will work together to manufacture, ship, and install Post-K, which is set for launch between 2021 and 2022. Post-K will be installed in the RIKEN Center for Computational Science, where the K computer is currently located.

In addition, Fujitsu will productize a commercial supercomputer using technology created for Post-K, and plans to begin global sales in the second half of the 2019 financial year.

Fujitsu also pledges to work with open source communities; for example, the company will promote the Arm ecosystem, utilize open source software with Post-K, and apply the results obtained with the supercomputer.

INTEL WINS US$500M CONTRACT TO BUILD US EXASCALE COMPUTER

Intel Corporation and the US Department of Energy have inked a contract valued at over US$500 million to build the US' first supercomputer, with one exaFLOP of performance.

In a tight race to be the world's first to reach the exascale, the supercomputer, named Aurora, is scheduled to be delivered to the Argonne National Laboratory in 2021.

The foundation of the Aurora supercomputer will be new Intel technologies designed specifically for the convergence of AI and traditional high performance computing at extreme computing scale. These include a future generation of Intel® Xeon® Scalable processor and a future generation of Intel® Optane™ DC Persistent Memory. Intel's sub-contractor, Cray, will also contribute their next-generation 'Shasta' family of supercomputing systems.

"The convergence of AI and high performance computing is an enormous opportunity to address some of the world's biggest challenges and an important catalyst for economic opportunity," said Mr. Bob Swan, Intel's CEO.

Breakthrough research projects that could benefit from the supercomputer include extreme-scale cosmological simulations, new approaches for drug response prediction, and materials for the creation of more efficient organic solar cells.

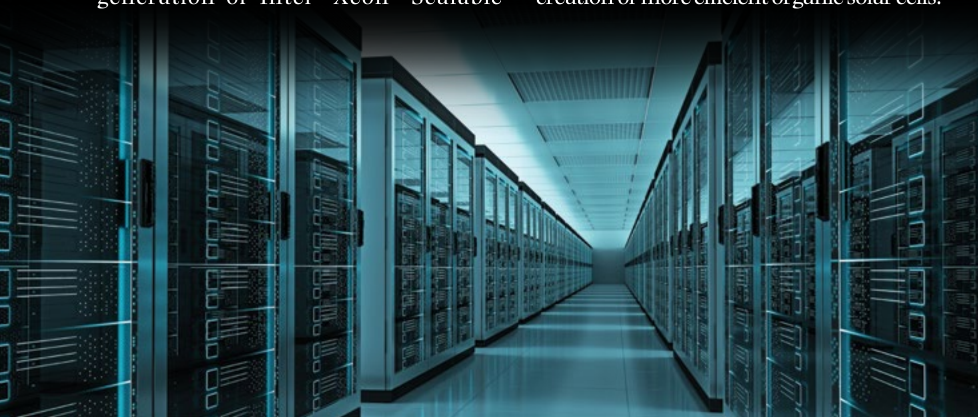

NVIDIA TO ACQUIRE MELLANOX FOR US$6.9B

NVIDIA announced that it has reached an agreement to acquire Israeli high-performance interconnect technology company Mellanox for approximately US$6.9 billion. This acquisition combines two of the world's leading companies in high performance computing that power over half of the world's TOP500 supercomputers.

The deal is expected to close at the end of 2019, with NVIDIA paying US$125 per share in cash in an ongoing consolidation in the world of processors, amidst interest from competitors such as Intel and Microsoft.

An early innovator in high-performance interconnect technology, Mellanox's InfiniBand and Ethernet interconnect technology are now used in over half of the world's fastest supercomputers and in many leading hyperscale datacenters. With Mellanox, NVIDIA aims to accelerate computing and build future datacenters that are able to handle the intensity of modern workloads in AI, scientific computing and data analytics.

"We're excited to unite NVIDIA's accelerated computing platform with Mellanox's world-renowned accelerated networking platform under one roof to create next-generation datacenter-scale computing solutions," said Mr. Jensen Huang, founder and CEO of NVIDIA.

SENSETIME & MALAYSIA'S G3 TO LAUNCH US$1B AI PARK

Malaysia Internet of Things (IoT) solutions provider G3 Global Bhd and China's SenseTime Group have announced plans to set up the first artificial intelligence park in Malaysia, in collaboration with China Harbour Engineering Company Ltd.

With a total investment of more than US$1 billion over five years, the park will serve as a platform for the development of AI solutions in areas such as computer vision, speech recognition, natural language, and humanoid and robotics.

In addition to the AI park, G3, through its owned subsidiary Atilze Digital Sdn Bhd, will facilitate the introduction of SenseTime's products into the Malaysian market. SenseTime, as one of the world's most prominent AI startups with a valuation of over US$4.5 billion, will provide industry expertise and knowledge in return.

Another objective of the partnership is to help grow Malaysia's fledgling AI industry. "The core idea is that the partnership is more than just taking products to sell. It's more for the long term—more of the sharing of knowledge," Mr. Md Radzi Din, CEO of Atilze, told *Digital News Asia*.

MALAYSIA TO HOST BILLION-RINGGIT GREEN DATA CENTER

A RM1.2 billion (~US$287 million) next-generation green data center and supercomputing facility will be built near Labu, Negri Sembilan, in Malaysia.

Expected to be the largest facility in Southeast Asia, it will host 4,064 racks over 11,150 square meters of space. The facility is designed to have a reduced data carbon footprint by up to 15 percent.

Spearheaded by Japanese-owned company Regal Orion Sdn Bhd, the data center is expected to commence operation by October 2020. It will host the mission-critical applications and data of Japanese companies in sectors such as information technology, telecommunication and automotives.

The 1,110 racks under phase one of the project have been fully rented out, and are expected to bring in a net profit of RM450 million (~US$10 million), said Mr. Ryuichi Shimokawa, chairman and owner of Regal Orion, to *The Star*.

With more Japanese companies looking to host their data outside Japan, Shimokawa said that Malaysia would be a perfect place for such facilities. "The country has good and sufficient electricity supply, is economically and geographically stable, and offers cheaper operating cost compared to Singapore and Japan," he said.

Super Snapshot

INDIA'S NEED FOR SPEED

Building India's data highway

India is determined to catapult itself into the league of supercomputing powerhouses currently dominated by countries such as China, Japan and the US. To that end, it announced in March 2015 a seven-year project called the National Supercomputing Mission, allocating it a budget of Rs 4,500 crore (~US$640 million).

KEY GOALS OF THE NATIONAL SUPERCOMPUTING MISSION:

- Train 20,000 high performance computing (HPC) professionals
- Connect the supercomputers via a high-speed pan-India network called the National Knowledge Network
- Create a cluster of 73 HPC centers, including three petascale supercomputers

THE FIRST THREE SUPERCOMPUTERS TO BE BUILT UNDER THIS PLAN:

1. **Indian Institute of Technology (IIT), Banaras Hindu University Param Shivay** 833 teraFLOPS
2. **IIT, Kharagpur** 1.3 petaFLOPS
3. **Indian Institute of Science Education and Research, Pune** 650 teraFLOPS

APPLICATIONS

- Climate modeling
- Weather forecast and natural disaster prediction
- Discovering new drugs and materials
- Aerospace engineering

FUN FACT: The US initially embargoed the export of Cray supercomputers to India due to concerns that they would be deployed for military use. Despite this roadblock, India assembled its first indigenous supercomputer, called PARAM 8000, in 1991. Since then, the peak performance of India's top supercomputer, Pratyush, has improved by more than six orders of magnitude, with a peak power of 6.8 petaFLOPS.

SCAsia 2020
Supercomputing

Organised by the National Supercomputing Centre (NSCC) Singapore, SupercomputingAsia 2020 (SCA20) is an annual conference that encompasses an umbrella of notable supercomputing and allied events in Asia. SCA20 will be held from 24 to 27 February 2020. The key objective of the SupercomputingAsia conference is to promote a vibrant and relevant HPC ecosystem in Asia. Delegates will be able to gain access to visionary insights from thought leaders in academia and industry, optimum networking opportunities and the Supercomputing community in Asia.

Gathering the Best of HPC in Asia

www.sc-asia.org

24 to 27 February 2020
Singapore

SCA19 Highlights

Attended by more than **700** delegates from more than **26** countries

SCA19 **Guest of Honour, Mr Heng Swee Keat,** Minister for Finance and Chairman of the National Research Foundation Board

Get the latest HPC technology updates

Opening Keynote Speaker, Dr Christine Ouyang, Distinguished Engineer and Master Inventor, IBM

Optimal networking opportunities

WHAT'S IN SCA20?

INDUSTRY
HPC Technology Updates & Case Studies

SCIENTIFIC
Paper Presentations at Supercomputing Frontiers Asia

STRATEGIC SHOWCASE
Strategic verticals on the use of HPC

CO-LOCATED HPC EVENTS

TUTORIALS & EXHIBITIONS

Organised by

SG NSCC SINGAPORE
National Supercomputing Centre

For more information about the conference, please contact us at **SCA20@sc-asia.org.**

1 Fusionopolis Way, Connexis South Tower, #17-01, Singapore 138632 | www.nscc.sg | NSCCSG

Information is correct at time of printing.