AsianScientist
Issue 12
July 2022

SUPERCOMPUTING ASIA

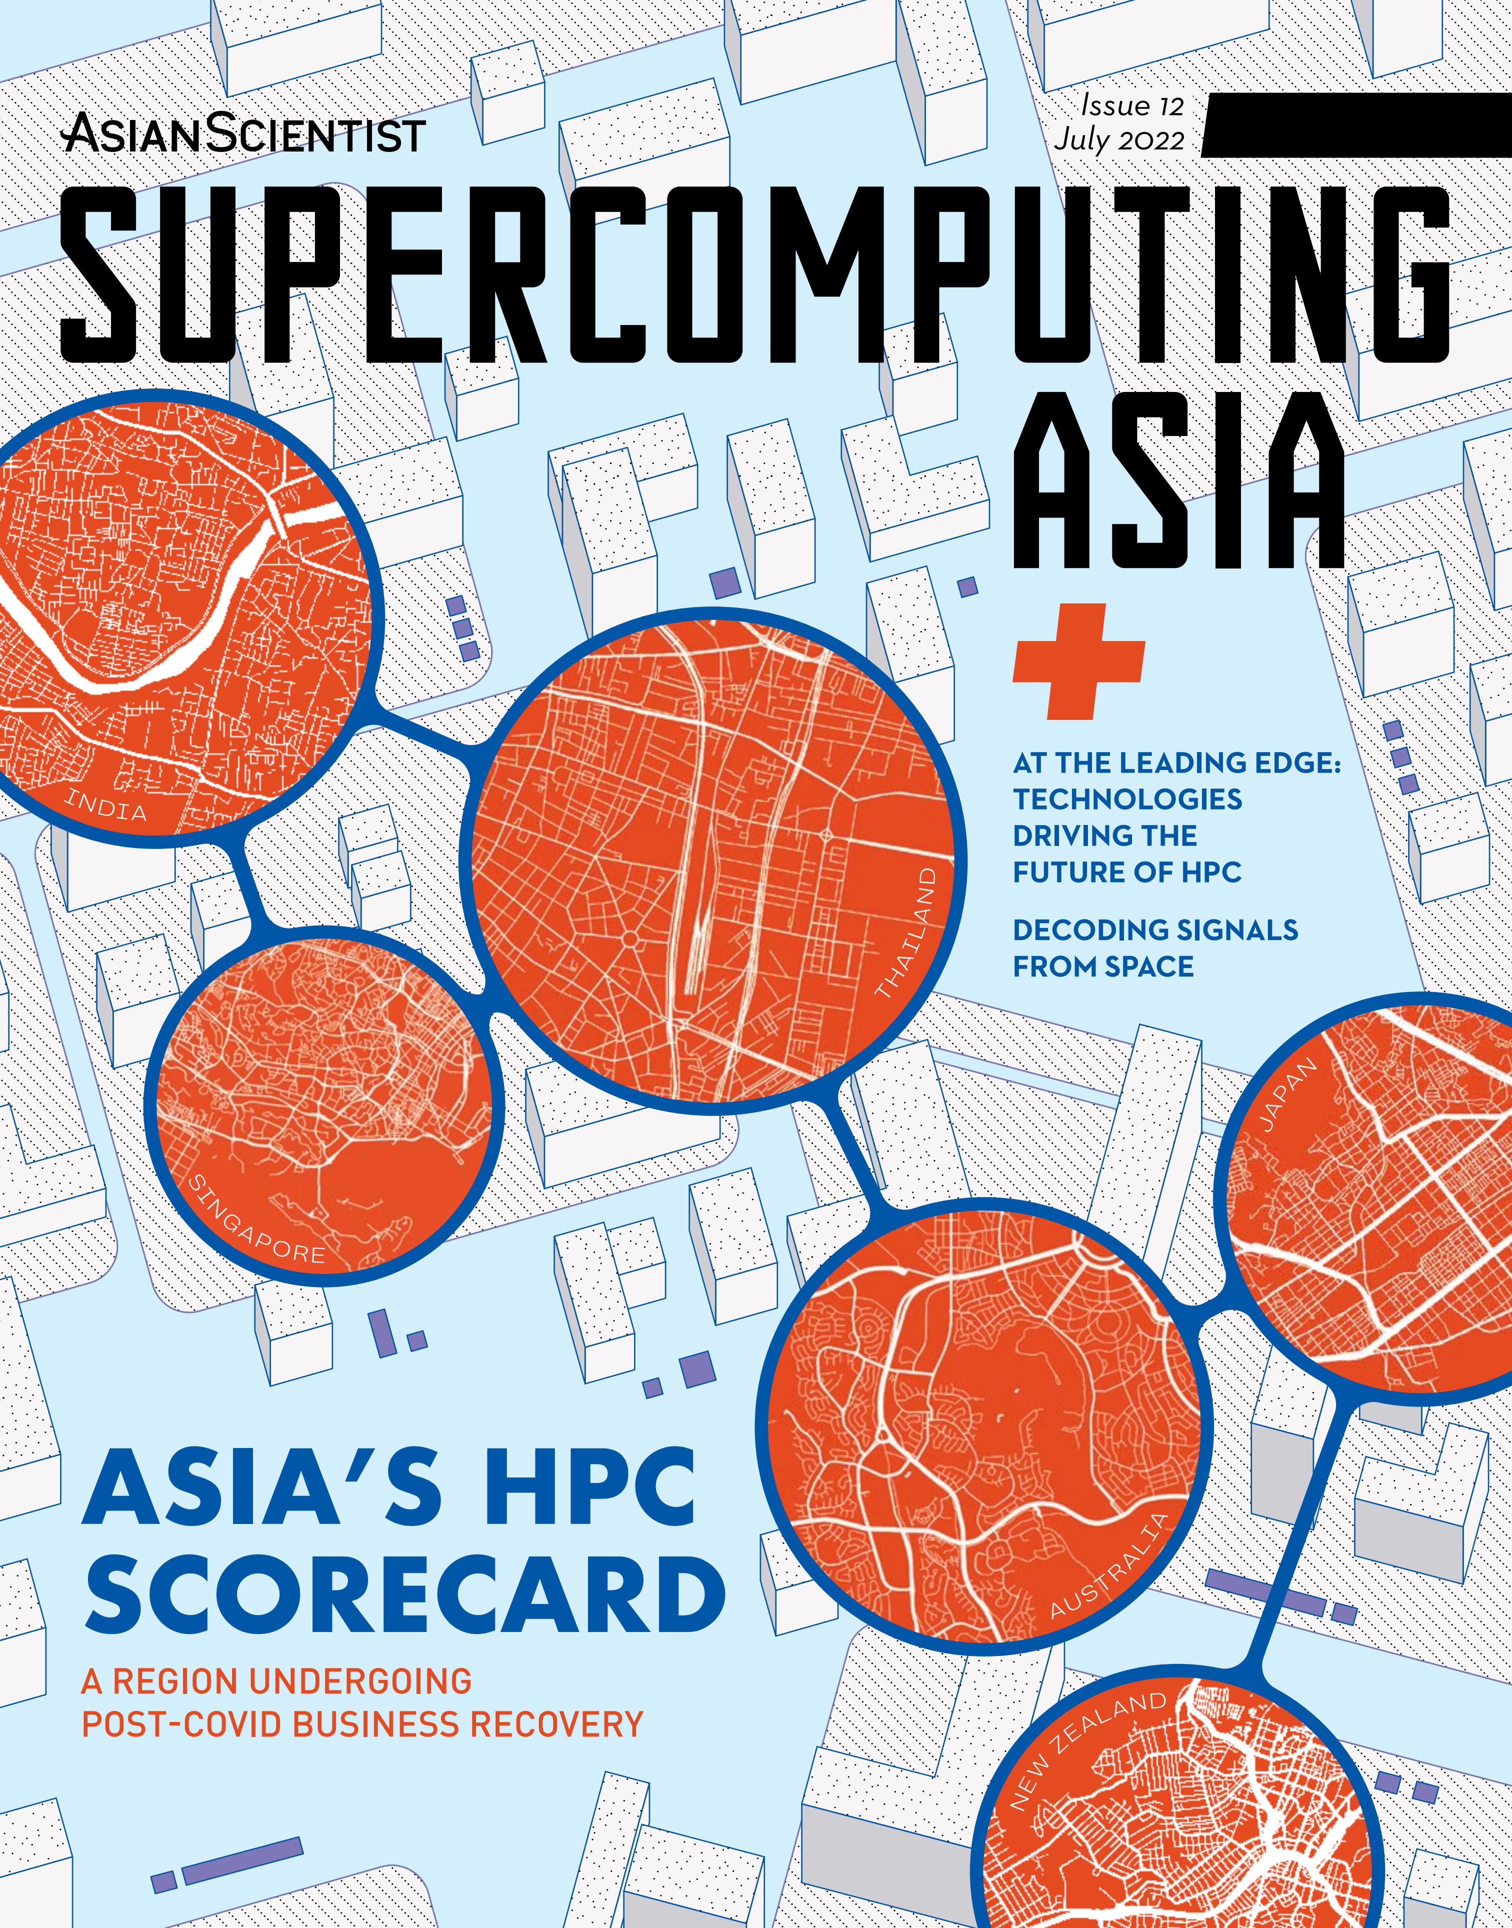

AT THE LEADING EDGE: TECHNOLOGIES DRIVING THE FUTURE OF HPC

DECODING SIGNALS FROM SPACE

ASIA'S HPC SCORECARD

A REGION UNDERGOING POST-COVID BUSINESS RECOVERY

CONTENTS

Issue 12
July 2022

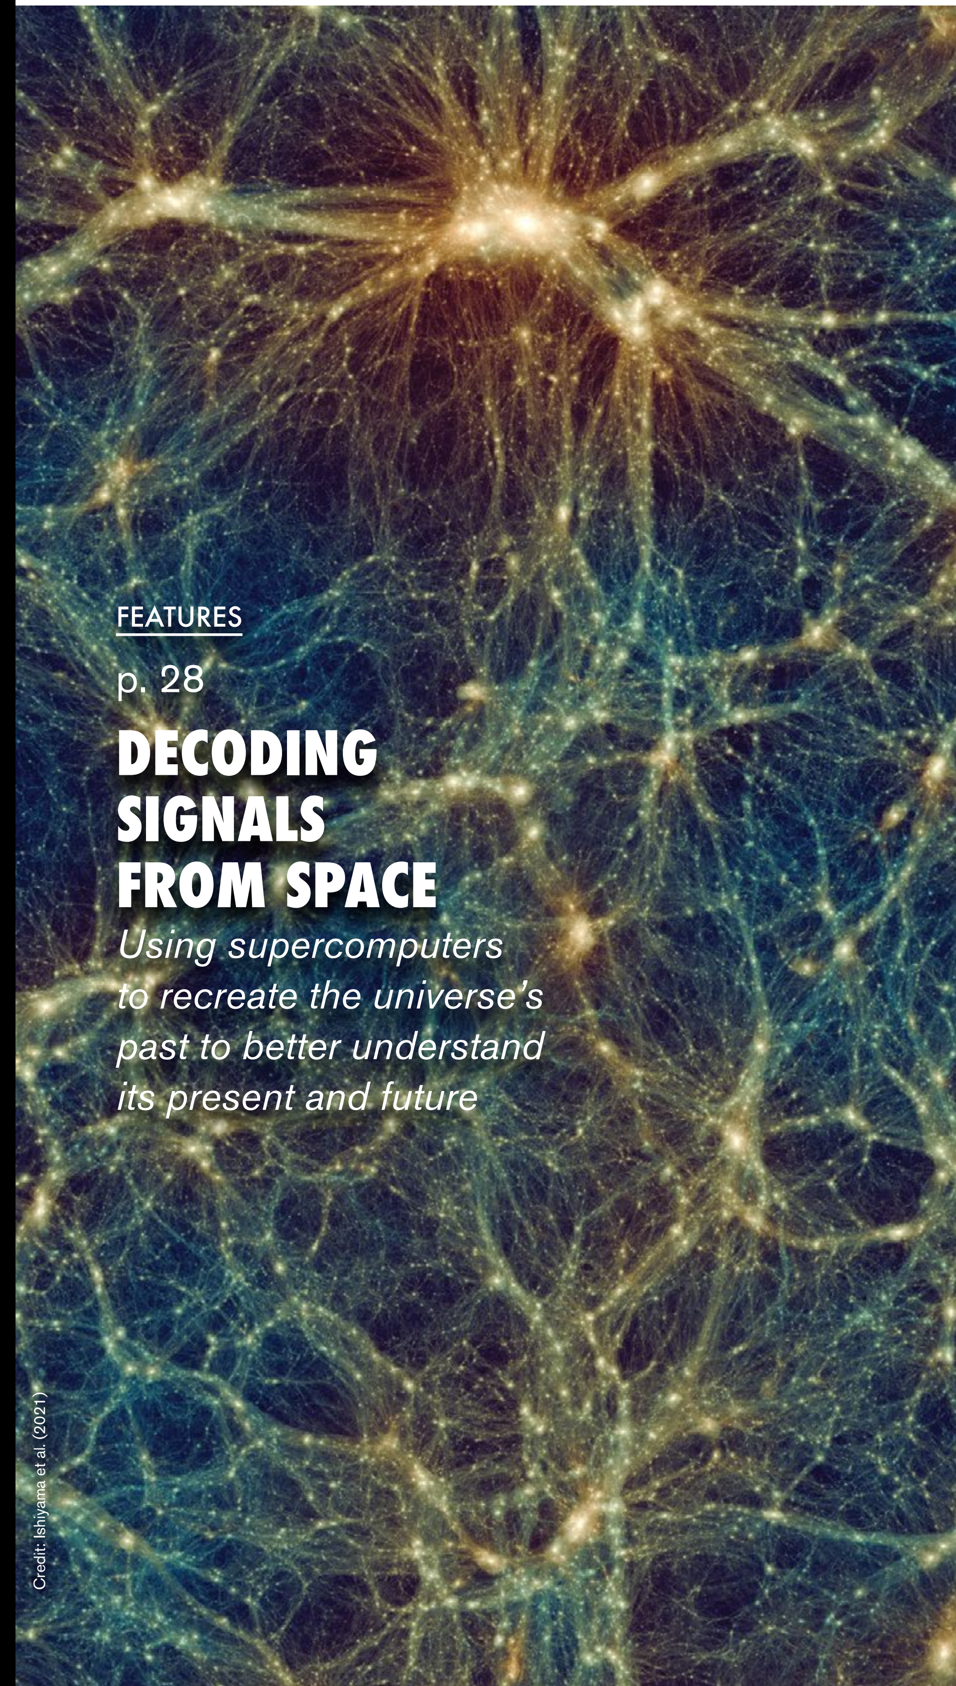

Credit: Ishiyama et al. (2021)

EDITOR'S NOTE

We are proud to enter the sixth year of *Supercomputing Asia* in print. The pandemic may be far from over, but we are optimistic about the role high performance computing will play in the COVID-19 recovery phase.

In the cover story (*Asia's HPC Scorecard*, p. 20), find out how the region is seeing increasing investment in supercomputers with their roles expanding rapidly across various sectors. In *Decoding Signals from Space*, (p. 28), read about how scientists from Australia and Japan are using supercomputers to reconstruct the universe's dynamic past to better understand its expansive present and future.

All over the world, leading supercomputing centers have begun to reach the lofty goal of exascale—those in Asia are no exception. In *Keeping Time on Asia's Race to Exascale* (p. 32), we showcase the potential towards sustained exascale performance of various supercomputers in the region.

Rounding out this issue, we share with you an infographic on Singapore's ASPIRE 2A supercomputer and its capabilities (*ASPIRE 2A in Numbers*, p. 38).

Enjoy the read and let's continue to celebrate technology!

Juliana Chan, Ph.D.
CEO & Publisher
Supercomputing Asia

/asianscientist
@asianscientist
Asian Scientist Magazine
asianscientist
AsianScientist
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

MANAGING EDITOR
Dr. Julian Tang

STAFF WRITERS
Jill Arul
Hannan Azmir

CONTRIBUTORS
Mitchell Lim
Erinne Ong

ART DIRECTOR
Shelly Liew

SENIOR DESIGNER
Ajun Chuah

DESIGNERS
Jorgina Tan
Yipei Lieu

SALES & MARKETING
Samantha Yeap
Jillene Kee
Yim Kong Joe

PUBLISHED BY
Wildtype Media Group Pte Ltd

DISTRIBUTED BY
Pansing Distribution Pte Ltd

HEAD OFFICE
Wildtype Media Group Pte Ltd
71 Ayer Rajah Crescent #06-26
Singapore 139951
hello@wildtype.media
Website: www.asianscientist.com/subscribe

WILDTYPE MEDIA

Combining savvy communication with technical rigor, Wildtype Media Group is Asia's leading STEM and healthcare media company, spanning digital, social media, video, 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.

www.wildtype.media

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

Digital Dispatch

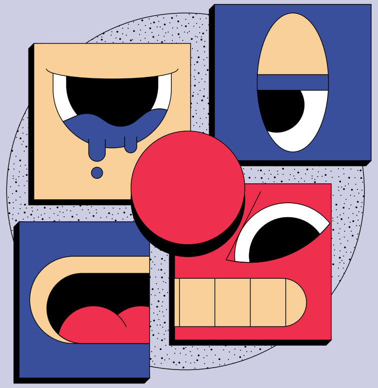

PREDICTING MENTAL HEALTH USING HIGH PERFORMANCE COMPUTING

At the start of the COVID-19 pandemic, researchers turned to high performance computing (HPC) and artificial intelligence to predict mutations in the SARS CoV-2 virus. However, decoding and predicting the development of mental illness is something that has yet to be attempted by computational medicine. A new collaboration in the US between Oak Ridge National Laboratory (ORNL), Cincinnati Children's Hospital, the University of Cincinnati and the University of Colorado is attempting to change that.

Using an algorithm on the country's largest supercomputer to rapidly parse through extremely complex data such as an individual's genetic traits, medical history and environmental exposure, the researchers hope to determine the trajectory of a child's mental health during their developmental years. These findings could lead to more accurate predictions of potential mental illnesses such as anxiety and could allow for early intervention and treatment before adulthood.

SATOSHI MATSUOKA: ACCOLADES FOR AN HPC PIONEER

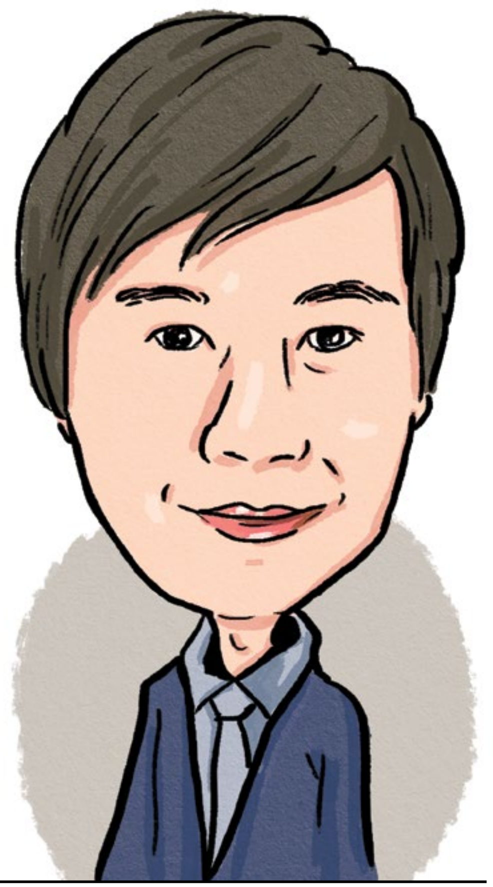

Professor Satoshi Matsuoka
Director of RIKEN's Center for Computational Science

Gaining recognition for their research is the ultimate dream for many researchers. For Professor Satoshi Matsuoka, this dream became a reality as he was awarded the Medal of Honor with Purple Ribbon by the Japanese government for his role in advancing HPC computing research in Japan.

Matsuoka is the director of RIKEN's Center for Computational Science, a member of the National Supercomputing Centre (NSCC) Singapore's steering committee and *Supercomputing Asia's* editorial board.

His most notable work was designing and leading the development of a cluster supercomputer series in Tokyo Institute of Technology's Global Scientific Information and Computing Center (GSIC), named the TSUBAME series. The current iteration of the TSUBAME series, TSUBAME3.0, was ranked first in the 2017 Green500 list of energy efficient supercomputers. He was also instrumental in the development of the Fugaku supercomputer, the world's top system for two years running, from 2020–2021.

Alongside this, he has also delved into the convergence of big data analysis, artificial intelligence and HPC. Some of his latest research include designing the next generation flagship supercomputer to succeed Fugaku, or 'FugakuNEXT', with innovative post-exascale features, for deployment towards the end of the decade.

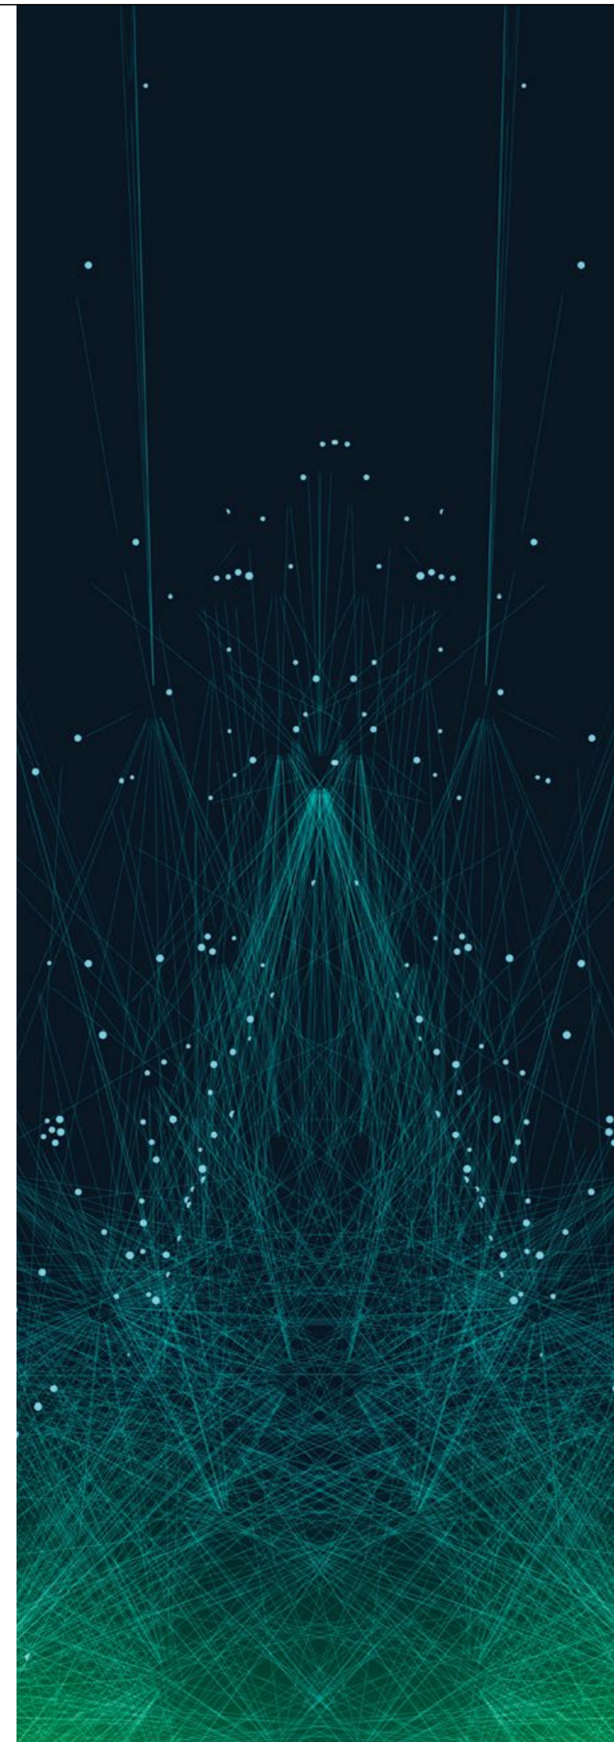

WINNING BIG IN MOVING DATA:

DATA MOVER CHALLENGE 2021 WINNERS ANNOUNCED

The SupercomputingAsia 2022 conference announced the winners of the International Data Movers Challenge 2021 (DMC 21). Organized by the National Supercomputing Centre (NSCC) Singapore, the competition brought together experts from industry and academia to showcase their innovative solutions for moving large data in today's interconnected age.

Seven teams from all over the world participated in this year's competition, with the focus being optimizing point-to-point wireless data transfers. Teams were given three months to develop and deploy their software, with a panel of judges who would determine the winners of the competition.

Each of this year's winners showcased their unique approach towards the challenge, remarked Andrew Howard, associate director of cloud services at the National Computational Infrastructure Australia and chief judge of the DMC 21 Judging Committee. From creating a new internet protocol to reliably send data across varying network conditions, to delivering large data sets through firewalls and impaired networks, each of the winners showcased their creativity and ingenuity in moving data.

Congratulations to all the winners:

Team MUSASHINO (Japan)
Most Innovative and Best IPv6 Performance

Arcitecta (Australia)
Most Complete Solution and Best Software Architecture

Fast Is Good – Raysync (Malaysia/China)
Best Virtualization Support

Globus (USA)
Best Integrated Software Experience

Ciena-iCair-UETN (Canada/USA)
Best Long-Distance Performance and Overall Winner

Digital Dispatch

ENHANCING AUSTRALIA-SINGAPORE HPC COLLABORATION

Australia's National Computational Infrastructure (NCI) and Singapore's National Supercomputing Centre (NSCC) have signed a memorandum of understanding for furthering advancements in supercomputing, particularly in areas such as the development of future HPC resources, education and training, and the exploration of more sustainable supercomputing operations.

Speaking to *PC World*, NCI's director Professor Sean Smith hopes that the new collaboration would further the interests of supercomputing, big data and high-throughput computing users in the Asia-Pacific region.

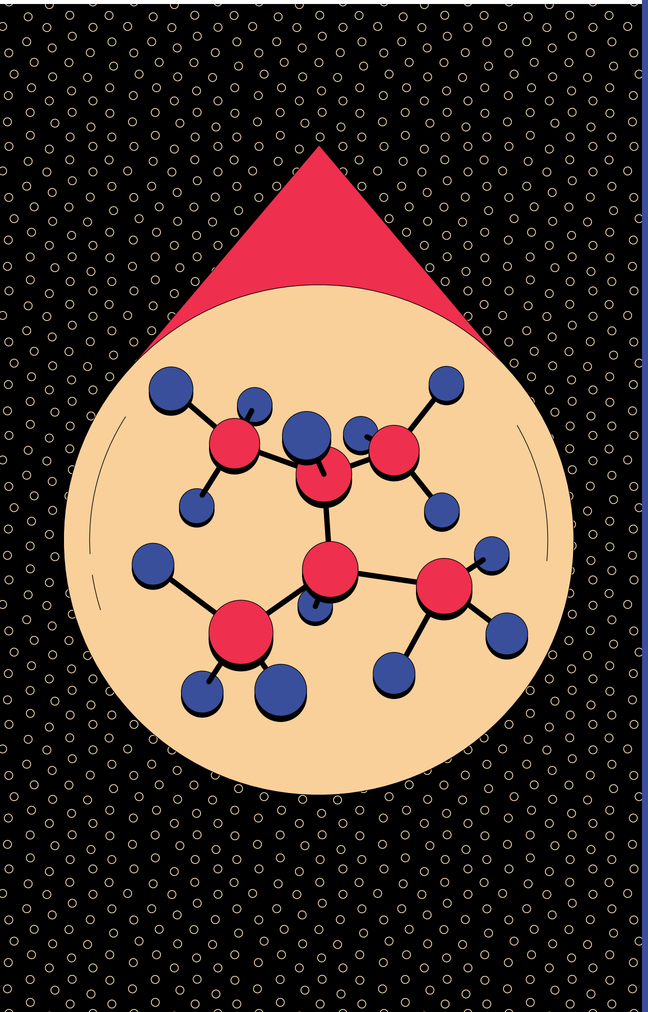

WHAT'S UP!

SC22

After last year's event attended by 3,500 in-person participants, the International Conference for High Performance Computing, Networking, Storage and Analysis 2022 (SC22) is getting ready for an even bigger event at the end of this year. The conference will be held in Dallas, Texas from November 13 to 18, with three days of workshops along with four days of exhibitions from the world's leading vendors, research organizations and universities.

With this year's theme being "HPC Accelerates", the key focus for the event will showcase the dizzying pace in recent breakthroughs from the HPC community. The rapid development of the COVID-19 pandemic along with the increased demand in faster data crunching and transfers created a huge boost in developing new computing systems, applications and software that keep pushing the limits of HPC.

SC22 will be showcased many successful uses of HPC—from predicting the weather in a world currently in the grips of climate change, to utilizing machine learning and large data analytics to advance the latest scientific discoveries—and participants will be engaged with leading industry experts, researchers and academics who have made waves in the HPC community.

Mark your calendars and get ready for an exhilarating event in high performance computing and beyond!

For more information, visit https://sc22.supercomputing.org/

WHERE

DALLAS, TEXAS, US

WHEN

NOV 13 – 18, 2022

HiPC 2022

The 29th edition of the IEEE International Conference on High Performance Computing, Data, Analytics, and Data Science (HiPC 2022) will take place in Bengaluru, India, on December 18 to 21, 2022.

HiPC 2022 serves as a forum for researchers around the world to present current work and activities on high performance computing systems, including data science and analytics, as well as their scientific, engineering and commercial applications. Discussions will range from traditional areas to emerging topics in the field.

For more information, visit https://hipc.org/

WHERE

BENGALURU, INDIA

WHEN

DEC 18 – 21, 2022

AT THE LEADING EDGE:

TECHNOLOGIES DRIVING THE FUTURE OF HPC

Take a look at the innovations that are unlocking the power of high performance computing to solve grand challenges in research and industry alike.

By **Erinne Ong**

Photo Illustrations by Shelly Liew / *Supercomputing Asia*

Much like the FIFA world rankings of the top football teams or the songs skipping up and down the music charts each week, the most powerful computers around the globe are also indexed in what is called the Top500 list.

For two straight years, Japan's Fugaku dominated the supercomputing charts, boasting a computing speed of 442 petaFLOPS. But a new challenger—the 1.1-exaFLOPS Frontier system at the Oak Ridge National Laboratory in the US—has made its debut atop the latest rankings released in May 2022, inching Fugaku down to the number two spot.

Besides the top places, the rest of the Top500 has also seen plenty of shuffling around during the list's biannual publication. Such movement in the rankings is a testament to the breakneck pace of technological advancement in the high performance computing (HPC) sector.

By providing high-speed calculations on vast amounts of data, HPC systems not only stand at the frontiers of the tech industry but also serve as enabling tools for tackling complex problems in many other fields. For example, scientists can use such technologies to uncover biomedical breakthroughs from clinical data or model the properties of novel materials more efficiently and accurately.

Given the ever-expanding value of these innovations, it comes as no surprise that researchers and industry leaders alike continue to challenge the ceiling for supercomputing—from components to clusters, minor tweaks to significant performance upgrades. As the promising potential of HPC relies on many moving parts, here are the technologies and trends that are laying the groundwork for building even more powerful and accessible supercomputing systems.

Revolutionizing machine intelligence

With the surge of data produced on a daily basis, artificial intelligence (AI) and data analytics tools are increasingly being used to extract relevant information and build models, which can then be used to guide decision-making or optimize systems. HPC is vital for enhancing AI technologies, including machine learning (ML) and deep learning (DL) systems built on neural networks that emulate the human brain's processing patterns.

Instead of analyzing data according to a predetermined set of rules, DL algorithms detect patterns and learn from a set of training data, and later apply those learned rules to new data or even to an entirely new problem. DL performance often depends on the amount and quality of data available—making it computationally expensive and time-consuming—but supercomputers can accelerate these learning phases and scour through more data to improve the resulting model.

In the medical sphere, for example, computational models simulate how intricate molecular networks interact to drive disease progression. Such discoveries can then spark novel ways to detect and treat complex disorders such as cancer and cardiometabolic conditions.

To investigate therapeutic targets against SARS-CoV-2, the virus that causes COVID-19, researchers from Chulalongkorn University in Thailand conducted molecular dynamics simulations using TARA, the 250-teraFLOPS supercomputing cluster housed at National Science and Technology Development Agency's Supercomputer Center. Through these simulations, the team mapped the interactions between a class of inhibitors and a protein known to be important for viral replication, generating new insights into how such drugs can be better designed to bind to the protein and potentially suppress SARS-CoV-2.

The power of HPC can also be harnessed for weather predictions and climate change monitoring, with South Korea building high-resolution and high-accuracy forecast models through its National Center for Meteorological Supercomputer. The Korea Meteorological Administration refreshed its HPC resources just last year to meet the extensive computational demands of climate modeling and AI analytics, installing Lenovo ThinkSystem SD650 V2 servers built on third-gen Intel Xeon Scalable Processors. Clocking in at 50 petaFLOPS, the new cluster is eight times faster and four times more energy efficient than its predecessor.

While supercomputing no doubt enables AI workloads, these smart systems can in turn be useful for optimizing HPC data centers, such as by evaluating network configurations for enhanced security. By monitoring server health, predictive algorithms can also alert users to potential equipment failures, helping reduce downtime and improve efficiency to support continuous HPC tasks.

A matrix of chips

HPC-powered AI may cover the software side of supercomputing, but the hardware is just as important. Advances in this space depend on innovations in developing processors or chips—pushing the boundaries of the number of operations that can be completed in as short a time frame as possible.

Perhaps the most familiar of these chips are the central processing units (CPUs), which can easily run simple models that process a relatively smaller volume of data. They typically have access to more memory and are designed to perform several smaller tasks simultaneously, making them useful for frequently repeated tasks but not for complex and lengthy work like training models.

Packing in more CPU nodes increases computing capacity, but just adding more units to the system is hardly efficient nor practical. To handle heavy ML workloads, accelerators in the form of graphical processing units (GPUs) and tensor processing units (TPUs) are critical to scaling up HPC resources—and in fact are the defining components that separate supercomputers from their lower-performing counterparts.

As the name suggests, GPUs excel at rendering graphics—no choppy videos or lagging frame rates in sight. But more than that, they are built to perform calculations in the nick of time, since smoothening out those geometric figures and transitions hinges on completing successive operations as quickly as possible. This speed enables GPUs to process larger models and perform data-intensive ML tasks.

TPUs push these computing capabilities a step further by taking care of matrix calculations more commonly found in neural networks for DL models than in graphical rendering. They are integrated circuits consisting of two units, each designed to run different types of operations. The unit for matrix multiplications uses a mixed precision format, shifting between 16 bits for the calculations and 32 bits for the results.

Operations run much faster on 16 bits and use up less memory, but keeping some parts of the model on 32 bits can help reduce errors upon executing the algorithm. With such an architecture, matrix calculations can be completed on just one TPU core rather than be spread out on multiple GPU nodes—leading to a significant boost in computing speed and power without sacrificing accuracy.

In the race to design better processors, chip manufacturing companies from all over the world are constantly exploring novel engineering methods and applying the latest research in materials science to elevate the performance of these critical HPC components.

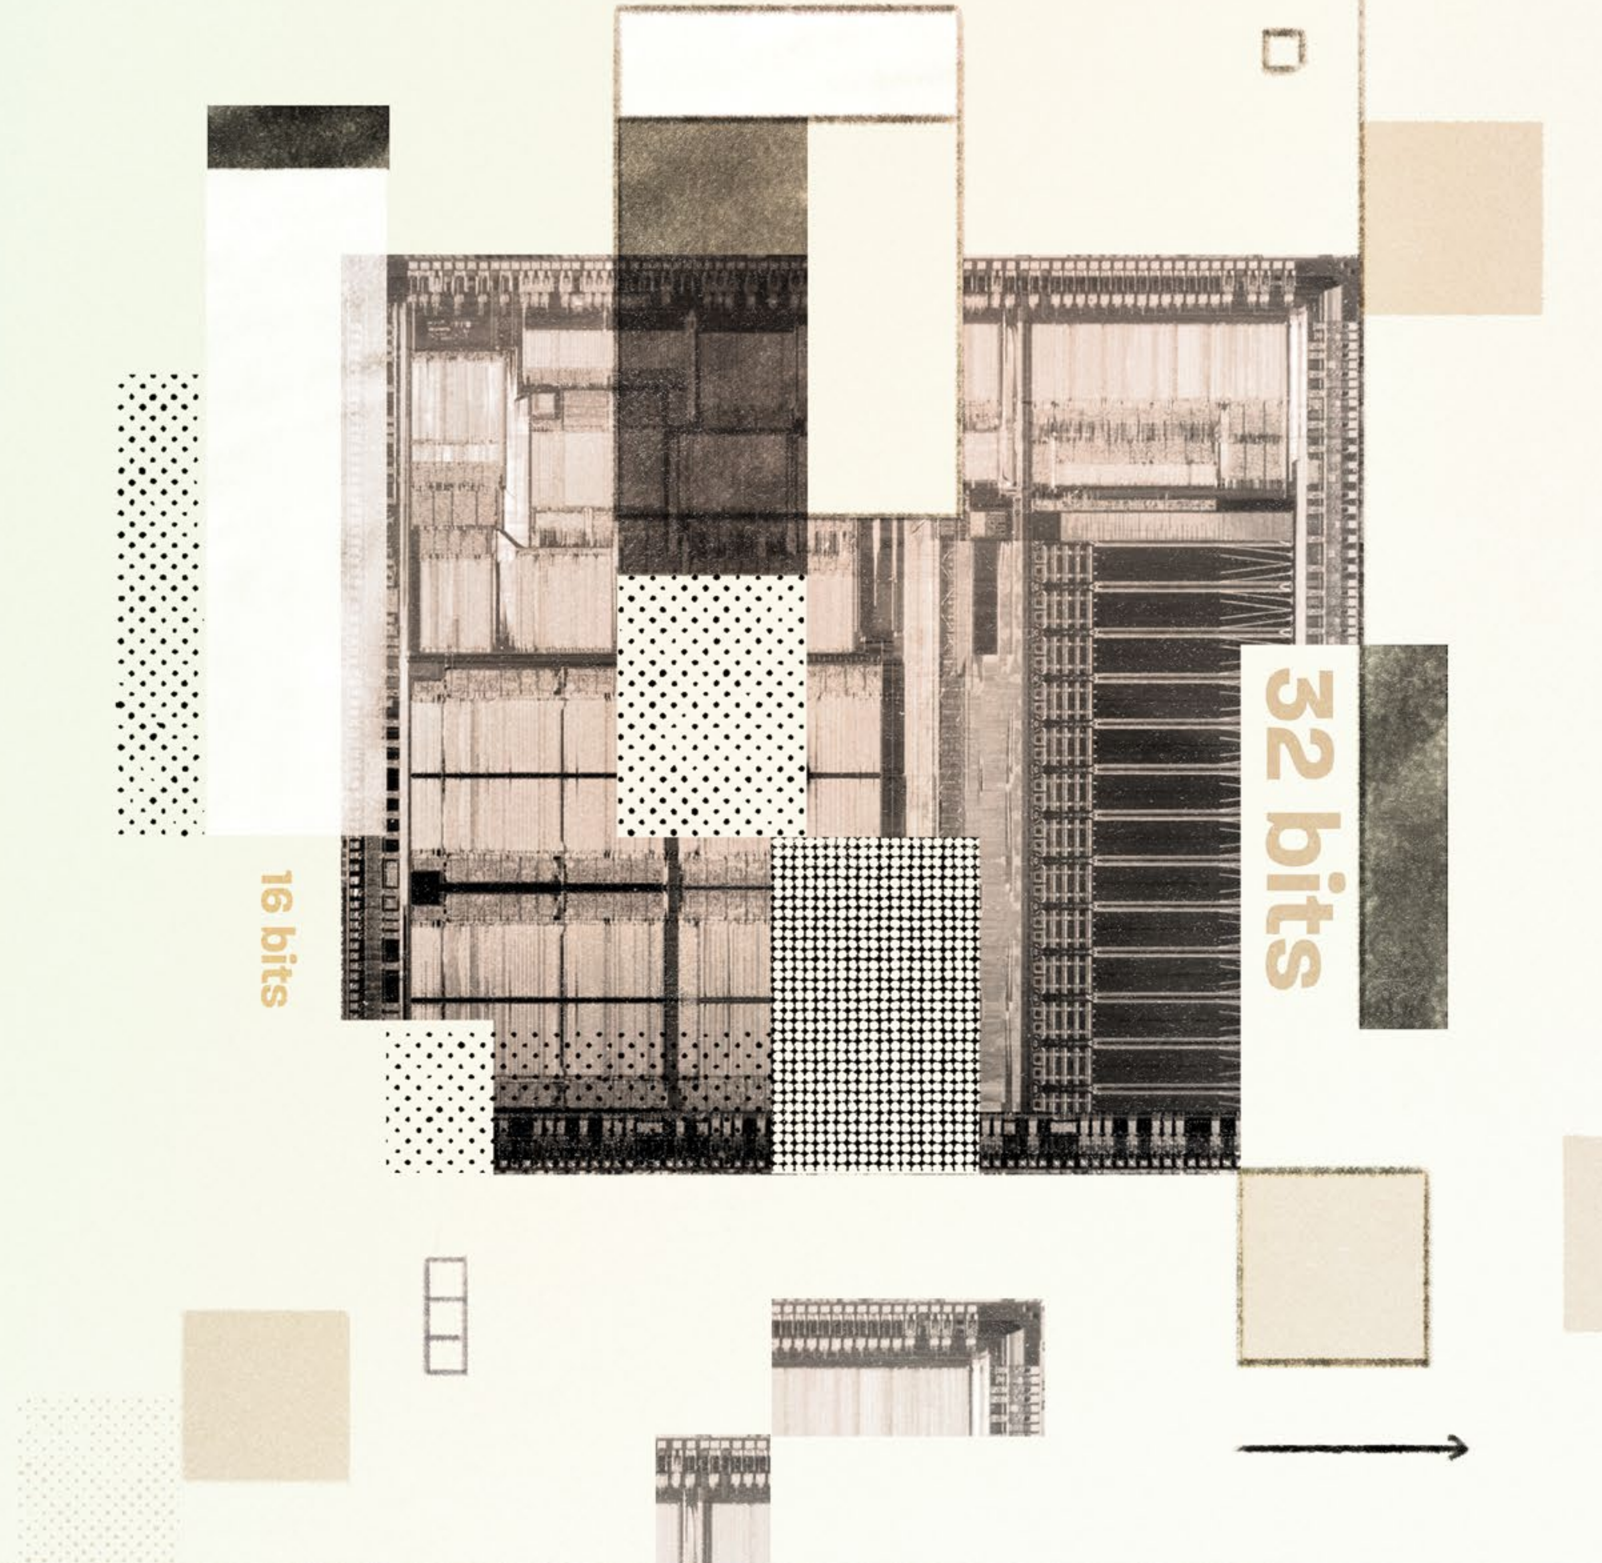

Accessing HPC resources on demand

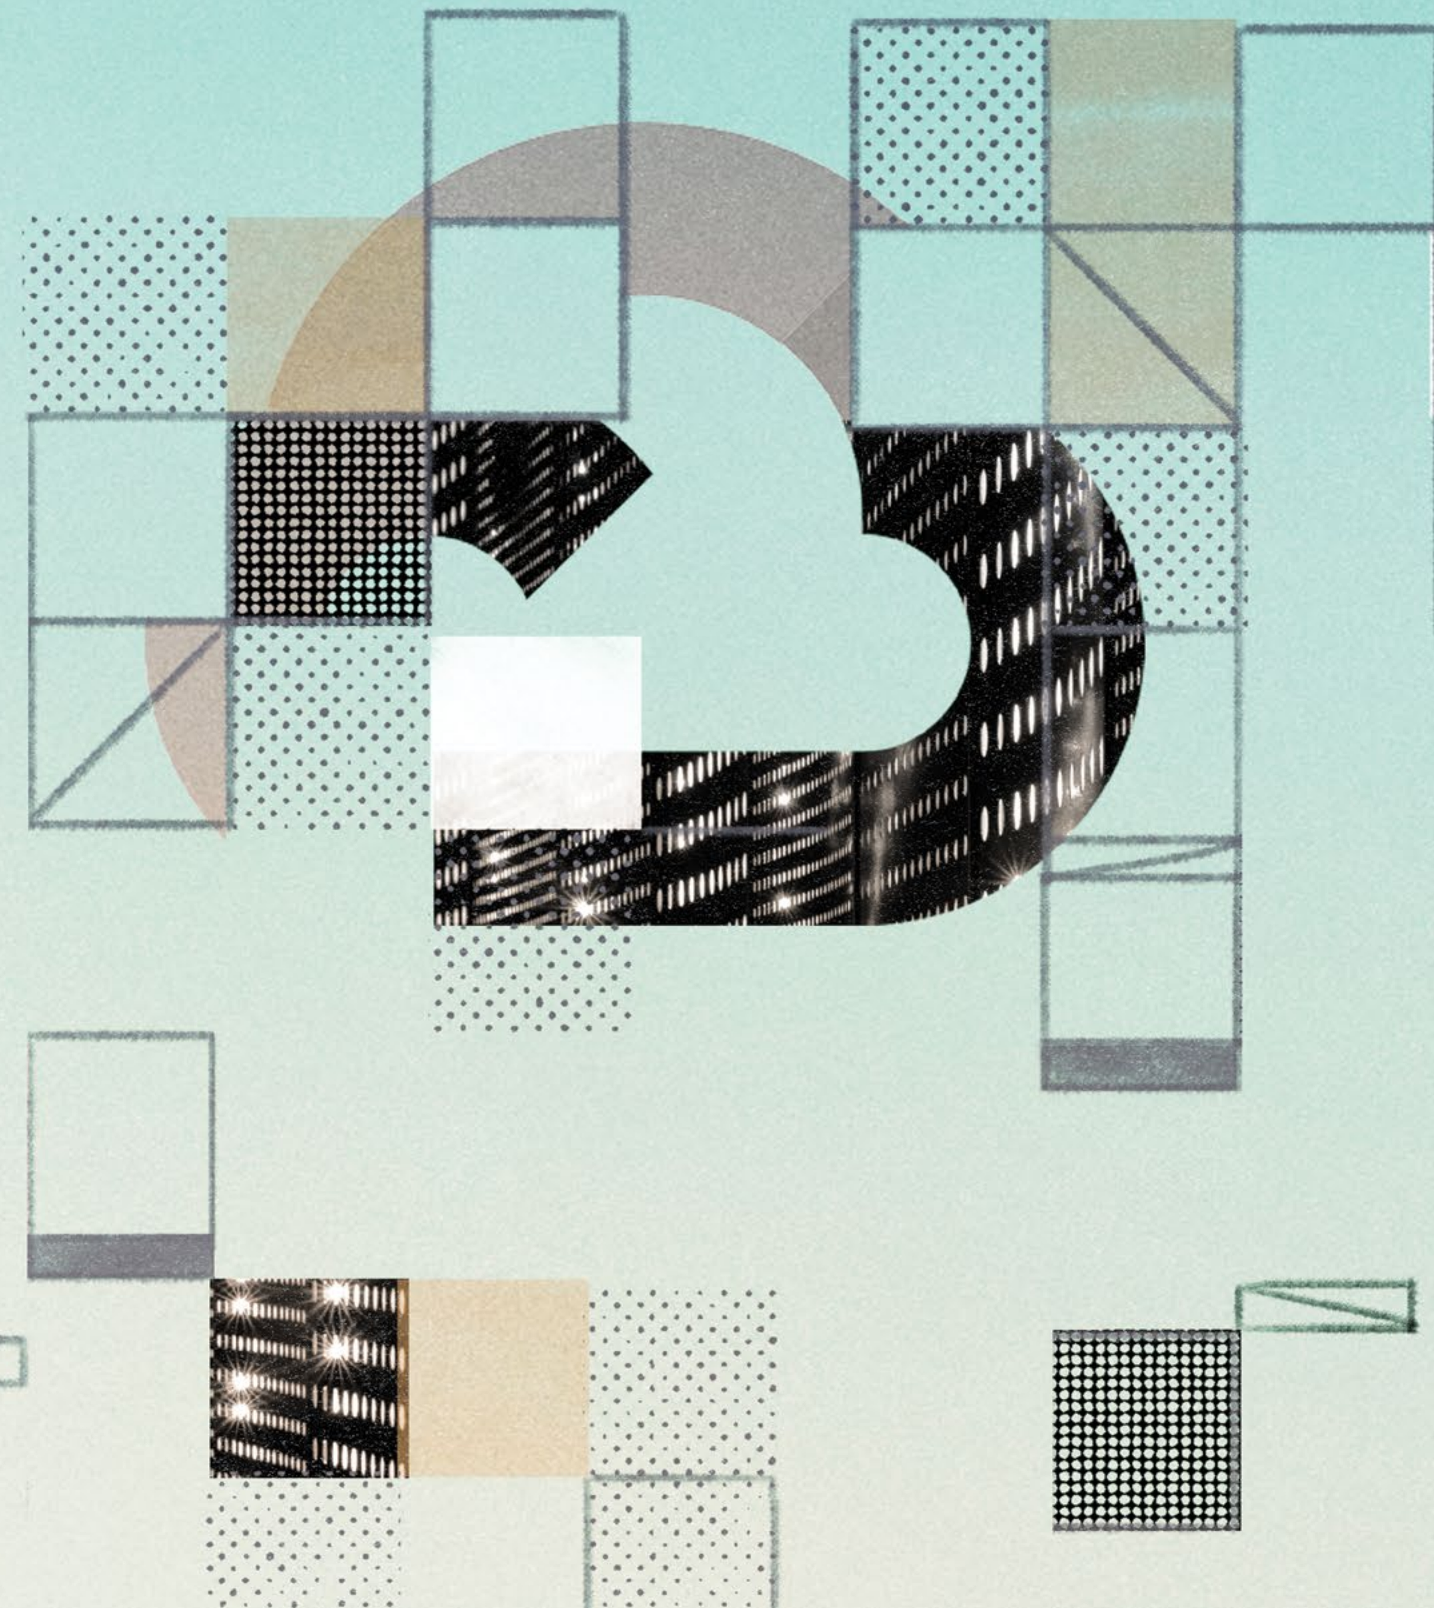

Supercomputing systems are hardly cheap—requiring significant financial, spatial and energy resources to build and maintain, not to mention the technical know-how to use them effectively. These costs can prove a barrier to widespread HPC adoption. Although HPC infrastructure is typically installed as in-house data centers, they have also been deployed on the cloud in recent years to increase access to these innovations.

Cloud computing involves delivering tech services over the internet, ranging from analytical processes to storage space. Called HPC as a Service (HPCaaS), this distribution of supercomputing resources across the cyberspace provides increased flexibility and scalability compared to on-site centers alone.

With supercomputing transitioning from academia to industry, HPCaaS can serve as an important bridge to place these powerful resources within the reach of more end users, from finance to oil and gas to automotive sectors. Through optimized scheduling strategies and allocation of resources, these systems can accommodate such diverse industry-specific workloads and encourage stronger collaborations over shared HPC capabilities.

In April this year, Japanese infocomms company Fujitsu—which jointly developed Fugaku alongside the RIKEN research institute—launched its HPCaaS portfolio with a vision to further spur technological disruption across industries. Through the cloud, commercial organizations can access the computational resources of Fujitsu's Supercomputer PRIMEHPC FX1000 servers, which run on ARM A64X processors and are supplemented by software for AI and ML workloads. These chips, which are also found in the Fugaku system, are not only high-end performers but are also very energy efficient.

To further encourage partnerships between academia and industry, Fujitsu is again working with the RIKEN research institute to ensure compatibility between the HPCaaS portfolio and the Fugaku system, granting more users and organizations the opportunity to use the region's most powerful supercomputer.

The HPC service's official release in the Japanese market is slated for October this year, and an international roll-out is also planned for the near future. By then, Fujitsu would also become the country's first-ever HPCaaS solutions provider, rivaling the infrastructure offerings of global companies including Google Cloud and IBM.

Flexible HPC consumption models will be key to bridging the digital divide, especially in Asia where technological progress is uneven and heterogeneous. By sharing top-notch resources, cross-border collaborations and the democratization of supercomputing can bring innovative ideas to life and carve new research directions with greater agility.

To the exascale and beyond

The arrival of Frontier marks an exciting milestone for the HPC community—breaking the exascale barrier. Prior to Frontier, the world's top supercomputers lived in the petascale when measured at 64-bit precision, with one petaFLOPS equivalent to a quadrillion (10^{15}) calculations per second.

These systems can execute extremely complex modeling and have advanced scientific discoveries at a swift pace. Fugaku, for example, has been used to map genetic data and predict treatment effectiveness for cancer patients; simulate the fluid dynamics of the atmosphere and oceans at higher resolutions; and develop a real-time prediction model for tsunami flooding.

Exascale computing could pave the way for even bigger breakthroughs, offering more realistic simulations and faster speeds at a quintillion calculations per second—that's 18 zeroes! This boost in speed can drive a diverse array of applications and fundamental research endeavors, such as understanding the complex physical and nuclear forces that shape how the universe works.

From sustainability to advanced manufacturing, scientists can also use these HPC resources to build more exact models of the Earth's water bodies, or dive deep into the nanoparticles and the optical and chemical properties of novel materials.

The chemical space is an especially exciting realm to explore, acting as the conceptual territory containing every possible chemical compound. Estimates are pegged at 10^{180} compounds—more than double the number of atoms inhabiting our universe, and a tantalizing figure relative to the 260 million substances documented so far in the American Chemical Society's CAS Registry.

Exascale computing can equip scientists with powerful new means to search every nook and cranny of this chemical space, whether for discovering potential drug molecules, light-absorbing compounds for solar cells or nanomaterials for more efficient water filters.

More compute resources can also support more distributed access and increased adoption of HPC, following in the footsteps of how the petascale systems were shared within and across borders.

While Asia may not yet have an exascale supercomputer on its soil, both Fugaku and China's Sunway have hit the exaFLOPS benchmark at 32 bits. With innovative minds at the forefront of the region's tech sector, achieving the same feat at the 64-bit level is on the horizon, boding well for the future of HPC and its applications in Asia and beyond.

Emerging from the global pandemic, the vitality and dynamism of Asia Pacific's supercomputing sector is almost palpable. From climate modeling to precision medicine, the frontiers of human knowledge are constantly being challenged—and this pursuit is not looking like it's losing momentum any time soon.

By **Mitchell Lim**

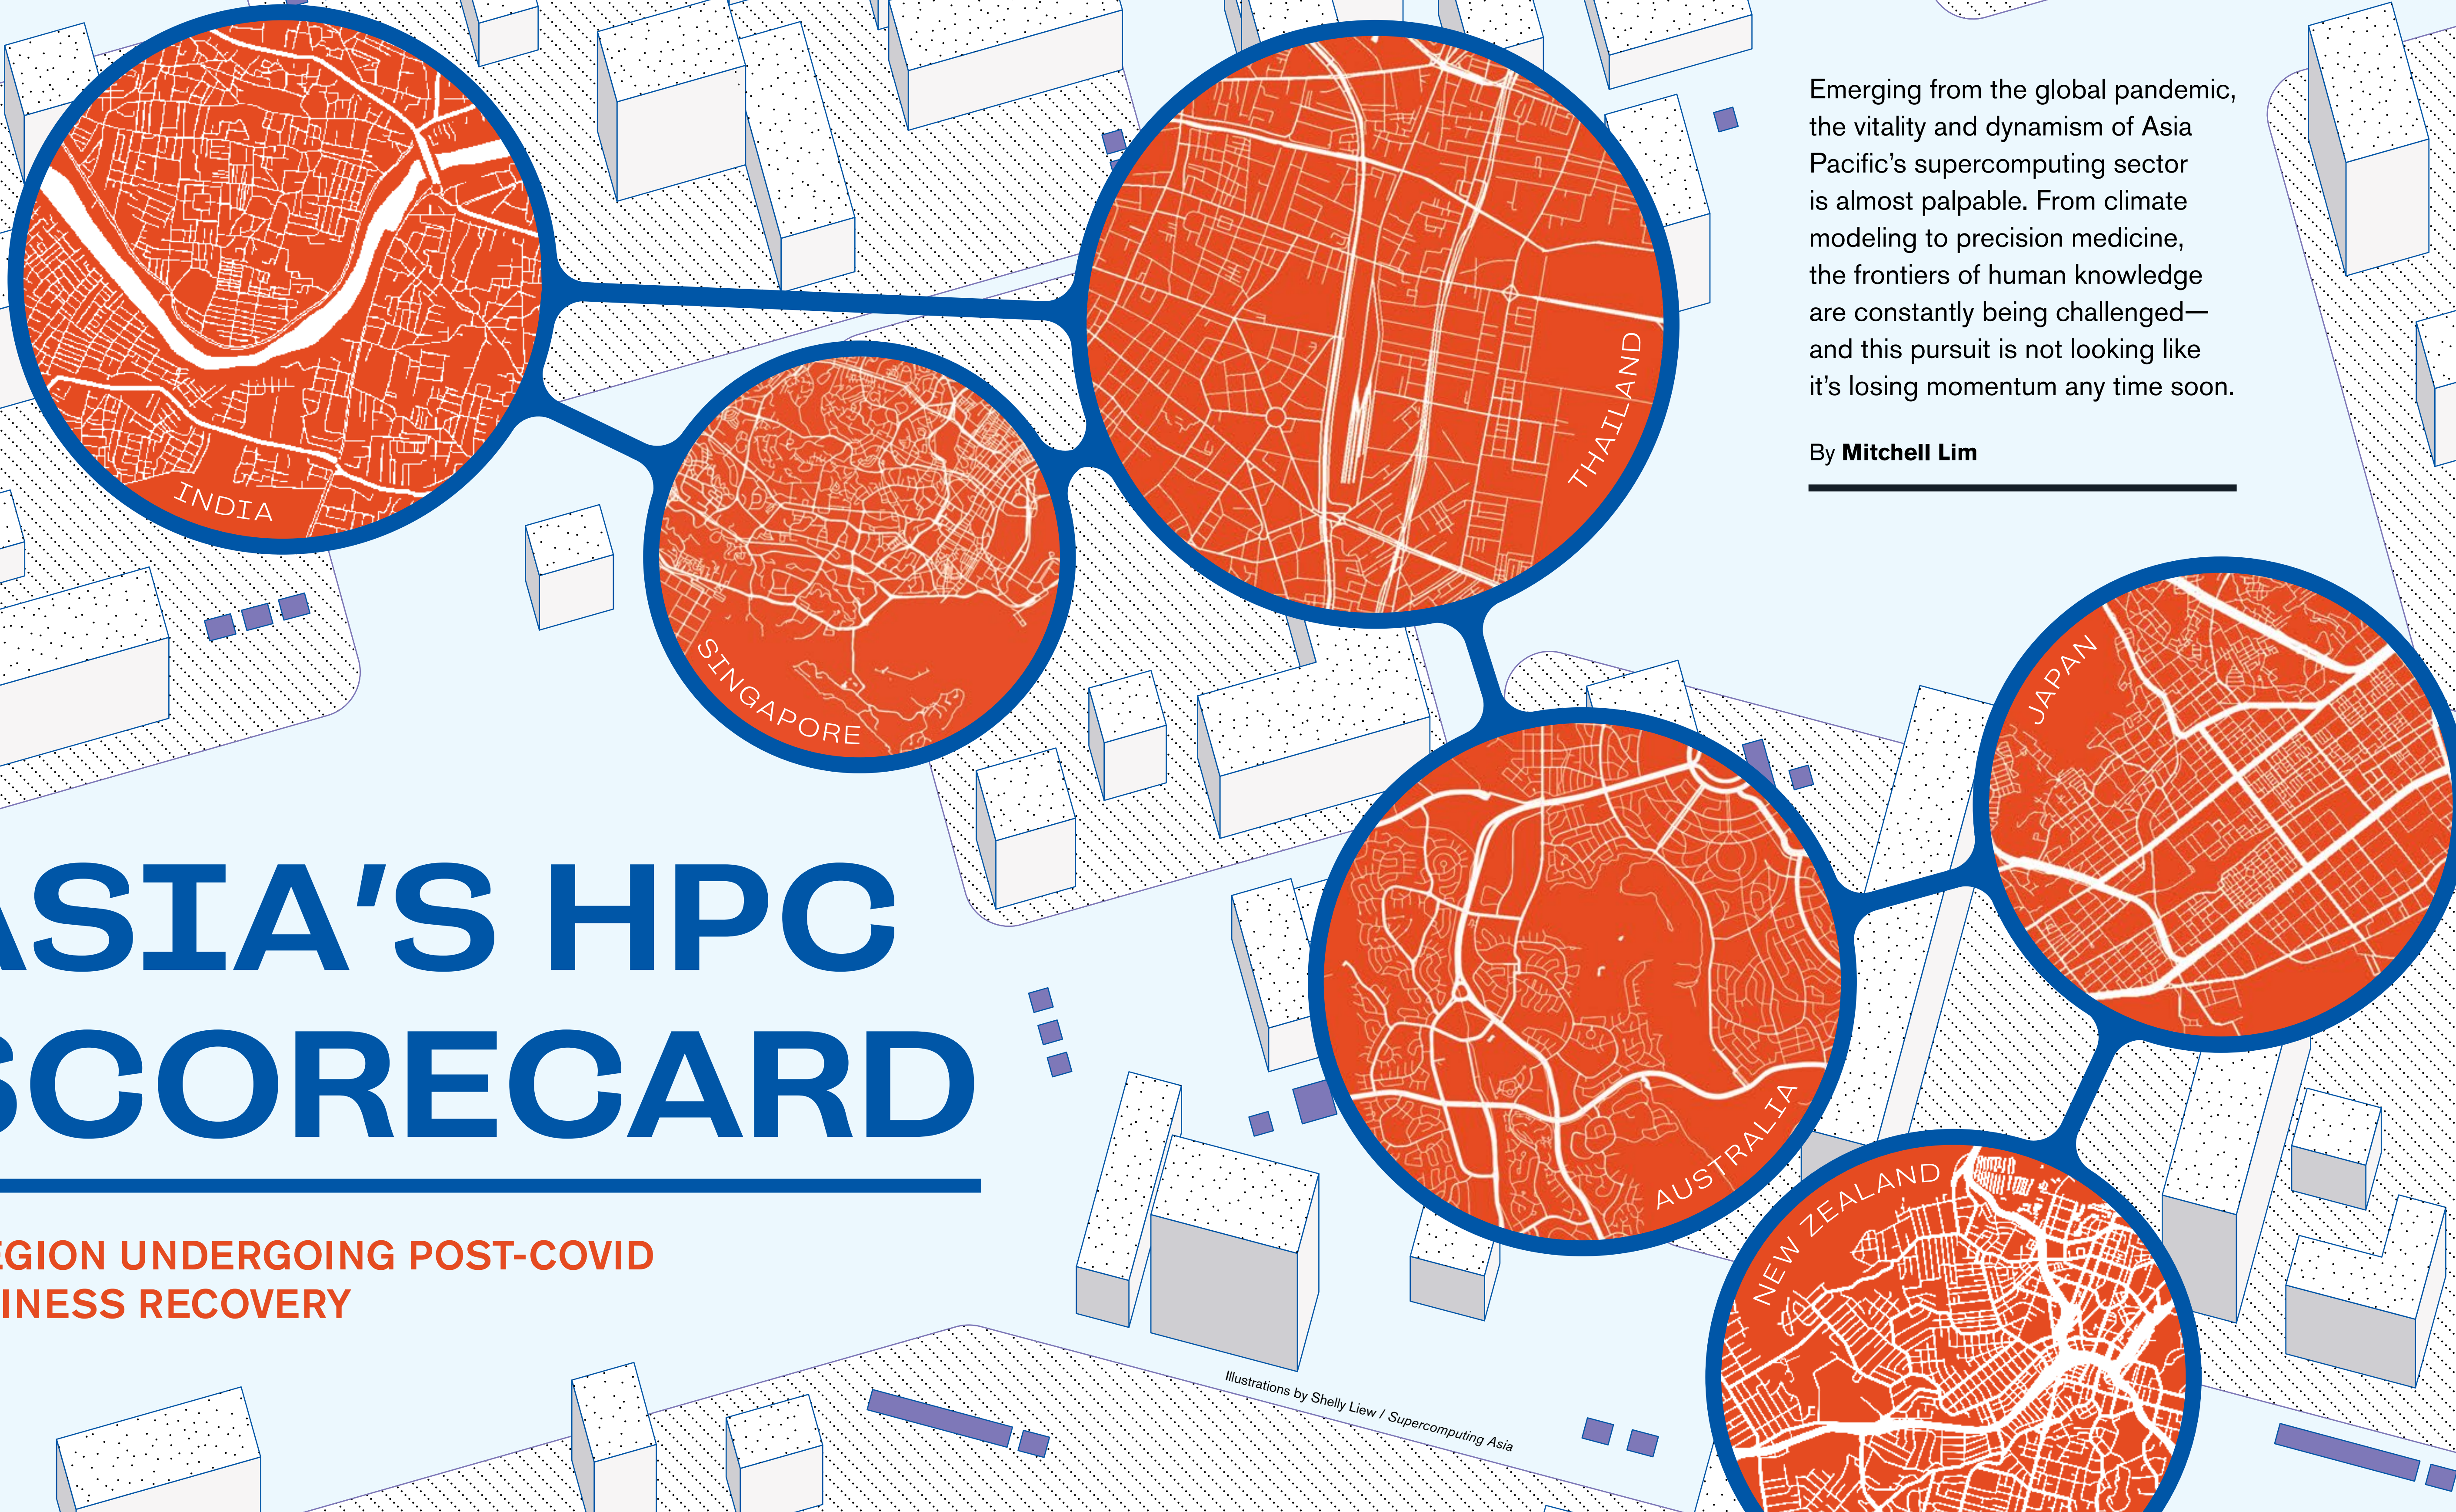

ASIA'S HPC SCORECARD

A REGION UNDERGOING POST-COVID BUSINESS RECOVERY

Illustrations by Shelly Liew / Supercomputing Asia

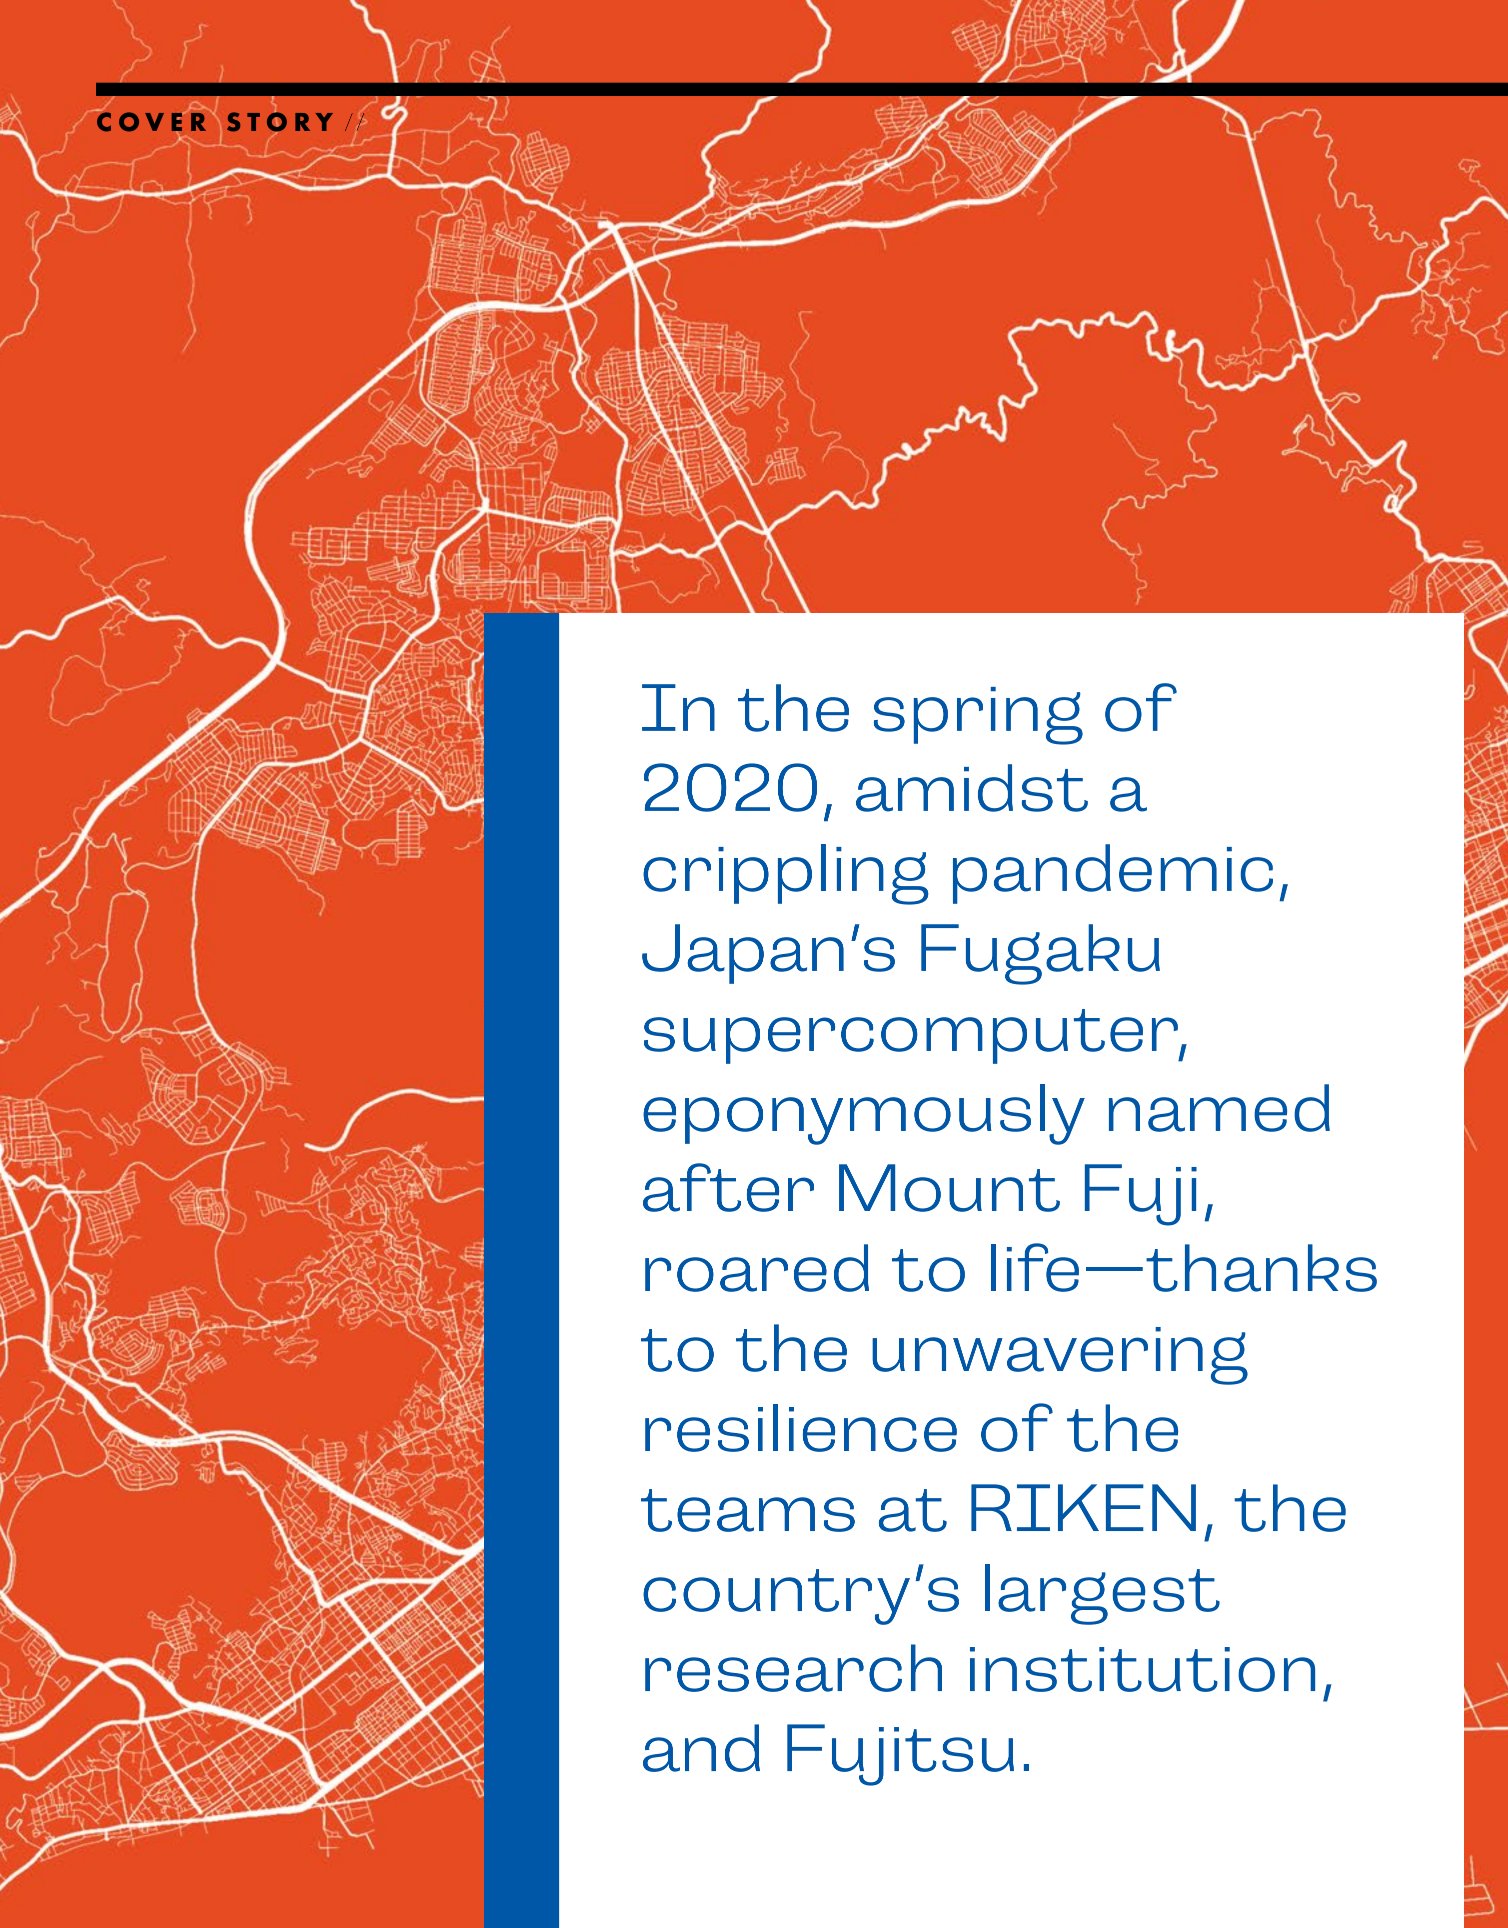

In the spring of 2020, amidst a crippling pandemic, Japan's Fugaku supercomputer, eponymously named after Mount Fuji, roared to life—thanks to the unwavering resilience of the teams at RIKEN, the country's largest research institution, and Fujitsu.

Named after the highest peak in Japan, Fugaku symbolizes the nation's aspirations for a supercomputer with towering capabilities. Indeed, shortly after its debut, the data-crunching beast scaled its way to the summit of the TOP500 list, making Asia home to the world's fastest supercomputer at the time.

"Fugaku as a supercomputer has never been about ranking," said Professor Satoshi Matsuoka, director of the RIKEN Center for Computational Science, during an interview with *Supercomputing Asia* in 2021. "Rather, our goal was to achieve application results and swiftly respond to important and difficult societal goals and problems."

STILL WELL IN ITS PRIME

Though Fugaku has recently lost its position as the world's fastest supercomputer to Frontier of the United States' Oak Ridge National Laboratory, it's noteworthy that Fugaku has held its crown for two consecutive years—an impressive feat that will no doubt leave behind a legacy.

Above all, Fugaku offered a beacon of hope in Japan's battle against the novel coronavirus. Made accessible almost a year ahead of schedule, the supercomputer turbocharged a swath of computationally intensive research to suppress the wrath of COVID-19. From deducing the recommended partition height to preventing virus spread, and from modeling the dispersion of droplets to assessing the effectiveness of face masks and more, Fugaku was instrumental in informing policy responses to keep the nation's infection rate at bay during the early stages of the pandemic.

As the world moves towards an endemic phase, Fugaku remains active in a panoply of scientific and industrial ventures. Marrying powerful compute with intelligent algorithms, researchers have revealed the origin story for carbon-12—an essential building block for life—in the turbulent cosmos, as well as slashing the runtimes of high-precision aircraft simulations from hours to mere minutes, among many other accomplishments.

While Japan is keen to build on Fugaku's achievements by grooming a successor, several countries in Asia are also the epicenter of exciting developments in high performance computing (HPC). Whether modeling molecular structures in quantum chemistry, assembling *de novo* genomes, evaluating the climate, or probing the depths of the universe, HPC is accelerating businesses and research workflows across research fields and industries.

"The massive amount of knowledge globally exceeds the ability of any individual clinician to master, and technologies such as supercomputers will play a vital role in processing or interpreting information to derive meaningful knowledge to advance medicine."

Associate Professor Yeo Khung Keong
Deputy Group Chief Medical Informatics Officer (Research) at SingHealth

HPC GIVING HEALTHCARE A BOOST

From artificial intelligence (AI) algorithms that read X-rays to personalized treatments underpinned by affordable genomic sequencing and population-wide health interventions impelled by big-data analytics, HPC has played a pivotal role in the evolution—and revolution—of the medical technology and healthcare landscape.

"Medical and health practitioners now have an increasing ease of access to data, which houses a vast range of information and resources," said Associate Professor Yeo Khung Keong, deputy group chief medical informatics officer (research) at SingHealth, in an interview with *Supercomputing Asia*. "The massive amount of knowledge globally exceeds the ability of any individual clinician to master, and technologies such as supercomputers will play a vital role in processing or interpreting information to derive meaningful knowledge to advance medicine."

Such a sentiment is also shared by the annual Supercomputing Asia conference held in March this year. The international conference, conducted in a hybrid format combining an in-person event with online components, emphasized the growing ubiquity of advanced supercomputing resources in a myriad of fields.

Among the many formal agreements made at the conference, a new Singapore-based public-private partnership aims to harness the power of supercomputers to solve complex healthcare issues. A tripartite collaborative effort, the project brings together the National Supercomputing Centre (NSCC) Singapore, Singapore's largest group of healthcare institutions SingHealth, and US-based chipmaker Nvidia Corporation. to develop and deploy a new supercomputer housed at the Singapore General Hospital.

As part of the project, Nvidia will breathe life into the compute engine by contributing their unique software development tools and pre-trained AI models. In addition, accessibility is on the project's radar—NSCC's Supercomputing Digital Sandbox environment will ensure researchers who don't speak the HPC language can still leverage the system for their work.

"Having supercomputing infrastructures equipped with AI and deep learning capabilities will provide unprecedented support for SingHealth's clinicians and researchers to better understand prevalent chronic diseases as well as gain insights into diagnosis and treatment," said Yeo.

In medicine, there's rarely a one-size-fits-all approach—two individuals with the same disease may respond very differently to the same treatment. This is where precision medicine comes in, treating patients based on individual genomic, phenotypic and lifestyle factors. In this regard, SingHealth seeks to further unlock the potential of the human genome with HPC. Using HPC-driven AI algorithms, medical scientists can better identify neoantigens, a repertoire of proteins that forms on mutating cancer cells that is specific to individual patients. Such discoveries will help to accelerate the development of potential personalized cancer therapies, which could directly destroy cancer cells with minimal to no side effects to the patient.

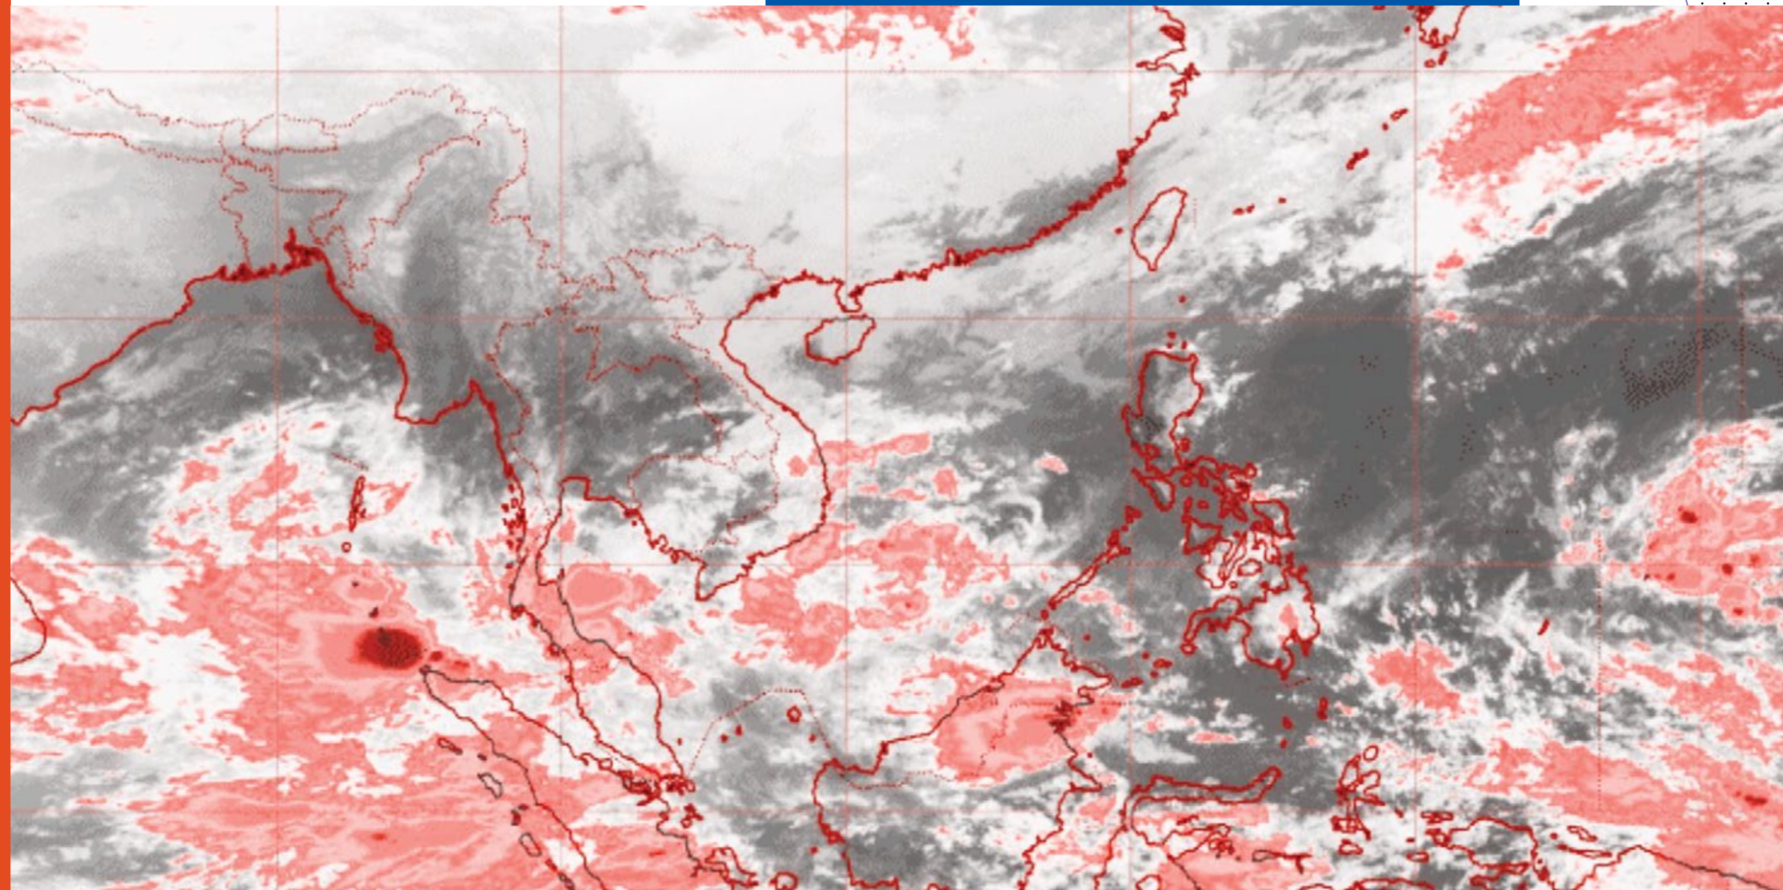

AN EMERGING SUPERCOMPUTING GIANT

Up north from Singapore, HPC has had its roots set in Thailand, dating back to the first Cray supercomputer installation in the early 1990s. Growing from its foundations, the country commonly known as the Land of Smiles is certainly offering more than just that—it has ambitious plans to expand its supercomputing network to take on new challenges that would drive progress across academia, government and industry.

Playing a vital role in realizing the country's HPC goals, the National Science and Technology Development Agency has teamed up with Nvidia and Hewlett Packard Enterprise (HPE) to build a new system at the Thai Supercomputer Center (ThaiSC). Touted as the largest of its kind in Southeast Asia, the new "LANTA" supercomputer, powered by HPE Cray EX, is adorned with a whopping 704 units of Nvidia A100 Tensor Core graphical processing units (GPUs)—twice as many as those in Singapore's latest ASPIRE 2A supercomputer, one of the fastest in the region.

"The system is created as an end-to-end, scalable HPC platform that supports a broad range of workloads," said Dennis Ang, senior director of enterprise business for ASEAN and ANZ at Nvidia, in an interview with *Supercomputing Asia*.

According to Dr. Anek Laothamatas, Thailand's Minister of Higher Education, Science, Research and Innovation, the new supercomputer is especially important for national disaster predictions. Being able to analyze vast quantities of climate data enables researchers to model and predict weather systems more accurately than what is achievable in a traditional set-up.

With adaptation now part and parcel of climate resilience, such research could arm industries from agriculture to architecture with panaceas in the face of adverse climate impacts. For instance, farmers can plant drought-resistant crops when the forecast calls for little to no rainfall. City planners can also adjust building codes for safer construction when destructive storms are anticipated.

"LANTA will provide HPC service on a national scale," said Dr. Piyawut Srichaikul, director of ThaiSC to *Supercomputing Asia*. "It will enable our scientists to break through the current constraints in computing resources and pursue frontier science in many competitive research areas."

"LANTA will provide HPC service on a national scale. It will enable our scientists to break through the current constraints in computing resources and pursue frontier science in many competitive research areas."

Dr. Piyawut Srichaikul
Director of the Thai Supercomputer Center (ThaiSC)

SHAPING THE FUTURE TOGETHER

Putting its pandemic woes behind it, India is set to become the world's fastest growing major economy for two consecutive years. On the HPC front, however, the nation has historically struggled to garner a supercomputing advantage—but that's about to change. The National Supercomputing Mission, announced in 2015, endeavors to catapult the nation as a frontrunner in HPC. Executed in three phases, involving assembly, manufacturing and design—all done locally—the mission will increase India's homegrown supercomputing capability to 45 petaFLOPS.

With ten supercomputers installed so far, this year sees nine more ready for action at leading Indian institutes. This will be a tremendous boon to hundreds of local institutes and thousands of active researchers working through the National Knowledge Network, the country's pillar for supercomputing systems.

> **"The aim of the roadmap is to encourage greater commercialization of research by allowing industry and researchers to engage more effectively."**
>
> **Melissa Price**
> Former Australian Minister for Science and Technology

REVAMP TO REINFORCE

Rather than reinventing the wheel, some strategic upgrades and updates could do the trick to refresh and reinvigorate a HPC ecosystem too. Outlined in the 2021 National Research Infrastructure Roadmap, a government strategy in Australia is set to provide the computing resources, digital tools, data governance frameworks and expertise needed to make best use of data. This includes establishing a more flexible, integrated computing ecosystem, where both national and commercial HPC resources are accessible to researchers.

"The aim of the roadmap is to encourage greater commercialization of research by allowing industry and researchers to engage more effectively," said former Australian Minister for Science and Technology Melissa Price in a press release.

Meanwhile, the New Zealand eScience Infrastructure (NeSI) is upgrading its Mahuika HPC cluster. Since it came online in 2018, Mahuika has witnessed its users increase 50 percent year-on-year, empowering the nation's scientists to tackle unique research ranging from simulating earthquake behaviors to managing biodiversity, and collaborating with the country's indigenous people and organizations to advance the wellbeing of New Zealanders.

A partnership between the University of Auckland, University of Otago and Manaaki Whenua—Landcare Research, the NZ$2.1 million investment will double the performance and boost the computational capabilities of the HPC system. This will ensure the island's national computing platforms remain responsive and high-performing to satisfy a growing array of data-centric and data-intensive research in biochemistry, physics, bioinformatics and more.

"We're seeing motivation within domains to advance and grow their own digital capability," said NeSI director Nick Jones, in an interview with *Supercomputing Asia*. "For example, there has been a dedicated drive within the genomics and bioinformatics fields to deliver and expand training and upskilling opportunities to its researchers."

NOT WITHOUT ITS CHALLENGES

Whether urgently addressing the climate crisis at hand, creating new medical innovations or fostering a more digitally woven society, the role of HPC is irrefutably central to navigating an unprecedented global transformation in many areas.

"Supercomputing in Asia has evolved along with the global HPC landscape," said Ang. "Collaboration is a key enabler—technology providers team up with local experts and agencies to ensure the latest innovations are available to power trailblazing research."

Nevertheless, advancing the frontiers of HPC in the region brings with it a set of equally challenging, and at times contradictory, problems that also require cross-border cooperation and collaboration.

"The unrelenting data explosion and the need for increasingly complex simulations call for even more computational power in the near term," added Yeo. "This means HPC will have a growing carbon footprint due to the associated increase in energy consumption, which is an obstacle that must be surmounted in the current state of the global climate."

"As a region we're recognizing common challenges in climate change and our environments, and we're inherently connected by our oceans," said Jones. "Therein lies our potential for greater connection and joint ambition. How we achieve this is open for exploration, and we're looking forward to learning more about how others are approaching these spaces."

> **"Therein lies our potential for greater connection and joint ambition. How we achieve this is open for exploration, and we're looking forward to learning more about how others are approaching these spaces."**
>
> **Nick Jones**
> Director of New Zealand eScience Infrastructure (NeSI)

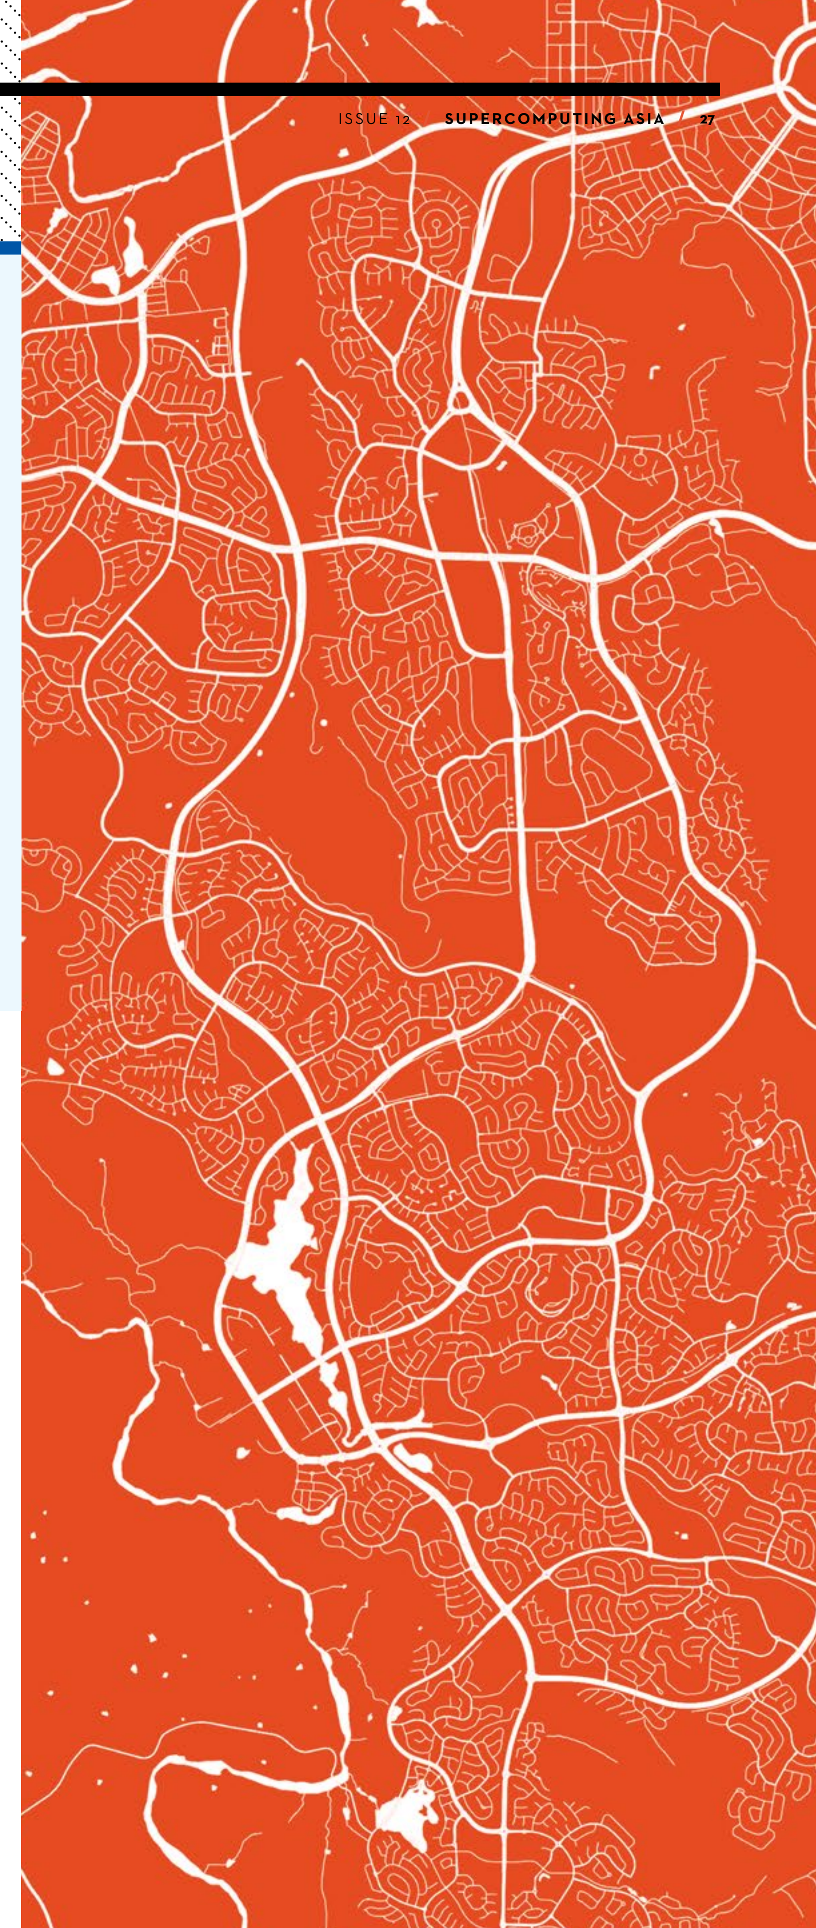

Decoding Signals From Space

Faced with datasets of galactic proportions, astronomers are turning to supercomputers to reconstruct the universe's dynamic past and better understand its expansive present and future.

By **Erinne Ong**

From hologram Princess Leia in the internationally beloved Star Wars franchise to Bubs the robot in the popular Korean movie *Space Sweepers*, which space-themed science fiction story would be complete without a matrix of futuristic tech? Aboard Starfleet vessels, an entire arsenal of computers—even artificial intelligence units—and handheld Personal Access Display Devices were all the rave in *Star Trek*.

But these fancy technologies are not just some pipe dream set in an unreachable future. In reality, supercomputing is part and parcel of how modern astronomy operates. Cosmology, a branch of astronomy dedicated to unraveling the origins and evolution of the universe, is particularly data-intensive and requires sophisticated computing resources to piece together disparate clues from outer space.

"The integration of astronomy and supercomputing has accelerated the rate at which discoveries can be made. We can process data much faster, detect far more faint signals due to leaps forward in sensitivity and create images at a higher resolution than ever before," said Dr. Sarah Pearce, deputy director for astronomy and space science at Australia's Commonwealth Scientific and Industrial Research Organisation (CSIRO), in an interview with *Supercomputing Asia*.

Pearce is also leading the charge for the Australian arm of the Square Kilometre Array (SKA) project, an international effort to build the world's largest radio telescope for collecting data over an area of a million square meters.

Technological juggernauts are critical to these astronomical missions, as one small step in high performance computing (HPC) can lead to one giant leap in understanding the cosmos and how our universe came to be.

GALACTIC ORIGINS

Twinkling stars and other space objects like asteroids are not just fascinating features that dot the sky, but they also hold many secrets about the fundamental forces in our universe, its immense history and its dynamic evolution. Even the seeming emptiness of outer space should fool no one, as invisible gravitational wavelengths and radio emissions fill the void with a colossal mishmash of signals.

"The sensitivity and design of the SKA telescopes will enable the detection of extremely faint signals emitted shortly after the birth of the universe almost 14 billion years ago," Pearce explained. "Like a time machine, such technologies allow us to look back to when and how the first stars and galaxies formed."

Origins aside, the universe is ever boggling the minds of astronomers. For one, it is still expanding—and doing so faster than ever before. The gravitational attraction among galaxies should slow down this expansion, yet a perplexing component called dark energy may be counteracting this force.

To test whether such theories hold weight, astronomers are taking stock of the masses of numerous galaxies and their gravitational disturbances on the path of radio waves. These galactic surveying missions also involve searching for hydrogen gas emissions, believed to fuel star formation.

Sampling an enormous number of galaxies is key to revealing the subtle differences in emission wavelengths and distortions in the radio signals. Accordingly, scientists are tapping into supercomputers to calibrate, transform and analyze all that data as quickly as possible—performing trillions of calculations in the blink of an eye. These measurements can then be used to build models to simulate the cosmological past.

For example, researchers led by Dr. Masato Shirasaki at the National Astronomical Observatory of Japan have turned back the cosmic clock and reconstructed the early universe, running 4,000 simulated universes on the 3.087-petaFLOPS ATERUI II supercomputer.

During the Big Bang, the universe exploded from nothing to one trillion-trillion times its size in a fraction of a second. This cosmic inflation influenced how galaxies and other matter are distributed in space. To retrace this phenomenon, the team stripped the simulated galaxies of their gravitational effects to reduce interference and evolved them to see which one best mirrored the state of the early universe.

"This new method allows us to verify inflation theories using only one-tenth of the amount of data," said Shirasaki to *Supercomputing Asia*. "Since less data is needed, it can also shorten the observing time required for future galaxy survey missions."

SKAVENGING FOR SIGNALS

To discover the answers to the universe's grand mysteries, scientists are devising machines that can match this galactic scale and decipher its cacophony of signals. Unlike their optical counterparts, radio telescopes like SKA can detect invisible waves and are not blocked by molecular dust, effectively peering into the "dark" regions where stars and planets are born.

The SKA low-frequency telescope in Western Australia is set to feature more than 130,000 antennas distributed across 512 stations, while the South African contingent will comprise 197 satellite dishes to cover the mid-frequency range.

"SKA will be receiving up to 10 billion data streams simultaneously," Pearce highlighted. "The supercomputers at our Science Processing Facilities will be integral to keep up with the data pouring in from the receivers 24/7."

Such extensive equipment can accelerate surveying missions by capturing several large parts of the sky in parallel and at unprecedented sensitivity. But to paint a picture out of the radio data, supercomputers need to correlate and synchronize the signals from the antennas, multiplying them together to generate data objects called visibilities.

"The difficulty is that inside these visibilities, the image of the sky is jumbled together with antenna responses and other radio signals such as from telecommunications devices," noted Pearce.

FROM SCOPES TO SUPERCOMPUTERS

Supercomputers employ advanced data analytics to disentangle space signals from all the noise, including accounting for minor differences in the instruments used and any "spikes" that appear around bright stars. Through iterative loops of calculations, the machines can convert the radio waves into astronomical images with unrivaled quality and resolution.

Whether filtering out interfering signals or stitching together smaller images to create detailed representations, these complex computing tasks all unfold in real time and are conducted over thousands of radio frequencies. Such a feat, Pearce noted, is possible only because of the sheer power of HPC resources available today.

"Distant galaxies, only glimpsed by very long observations today, will routinely be observed in a fraction of the time. Astronomers using the SKA telescopes will encounter more data than has ever been available in the history of radio astronomy," she added.

SKA also builds upon longstanding precursor projects from CSIRO, including the Australian Square Kilometre Array Pathfinder (ASKAP) and Murchison Widefield Array. The backbone of these space missions is Galaxy, a real-time supercomputing service for telescopes and astronomy research. Housed at the Pawsey Supercomputing Research Centre in Australia, this 200-teraFLOPS CRAY XC30 system is equipped with Nvidia K20X 'Kepler' graphical processing units and Intel Xeon E5-2690 host processors.

SKA's HPC facilities will boast a collective computing capacity of 500 petaFLOPS and archive over 600 petabytes of data each year. Moreover, the alliance of SKA centers across the globe will be connected through a high-end fiber network that can transmit data at speeds of seven to eight terabits per second—about 100,000 times faster than current average broadband rates.

> "Distant galaxies, only glimpsed by very long observations today, will routinely be observed in a fraction of the time."
>
> **Dr. Sarah Pearce**
> Deputy Director for Astronomy and Space Science at CSIRO

Credit: Ishiyama et al. (2021)

COLLECTIVE AMBITIONS, UNIVERSAL FUTURES

By bridging localized efforts to global endeavors, Pearce envisions a more collaborative model for the future of astronomy. "Deeply rooted in our ethos is the concept of open science. After a proprietary period, SKA's enormous data sets will become accessible for anyone who wants to analyze them, which dramatically boosts the potential for further discoveries."

Traditionally, a single astronomer or small team would request time to use a telescope for their individual research. Now, scientists and engineers from around 100 organizations across 20 countries are participating in the development of SKA—harnessing shared technological resources as the vehicle for driving advances in space science.

From revealing the secrets of dark matter to mapping the magnetic fields that permeate the universe, HPC systems are set to super-charge the next generation of astronomical observation. By capturing snapshots of space-time, these innovations can empower scientific teams to weave together riveting narratives that transform our understanding of the origin and fate of the universe.

KEEPING TIME ON ASIA'S RACE TO EXASCALE

All over the world, leading supercomputing centers have begun to reach the lofty goal of exascale computing.

By **Jill Arul and Hannan Azmir**

Illustrations by Shelly Liew / *Supercomputing Asia*

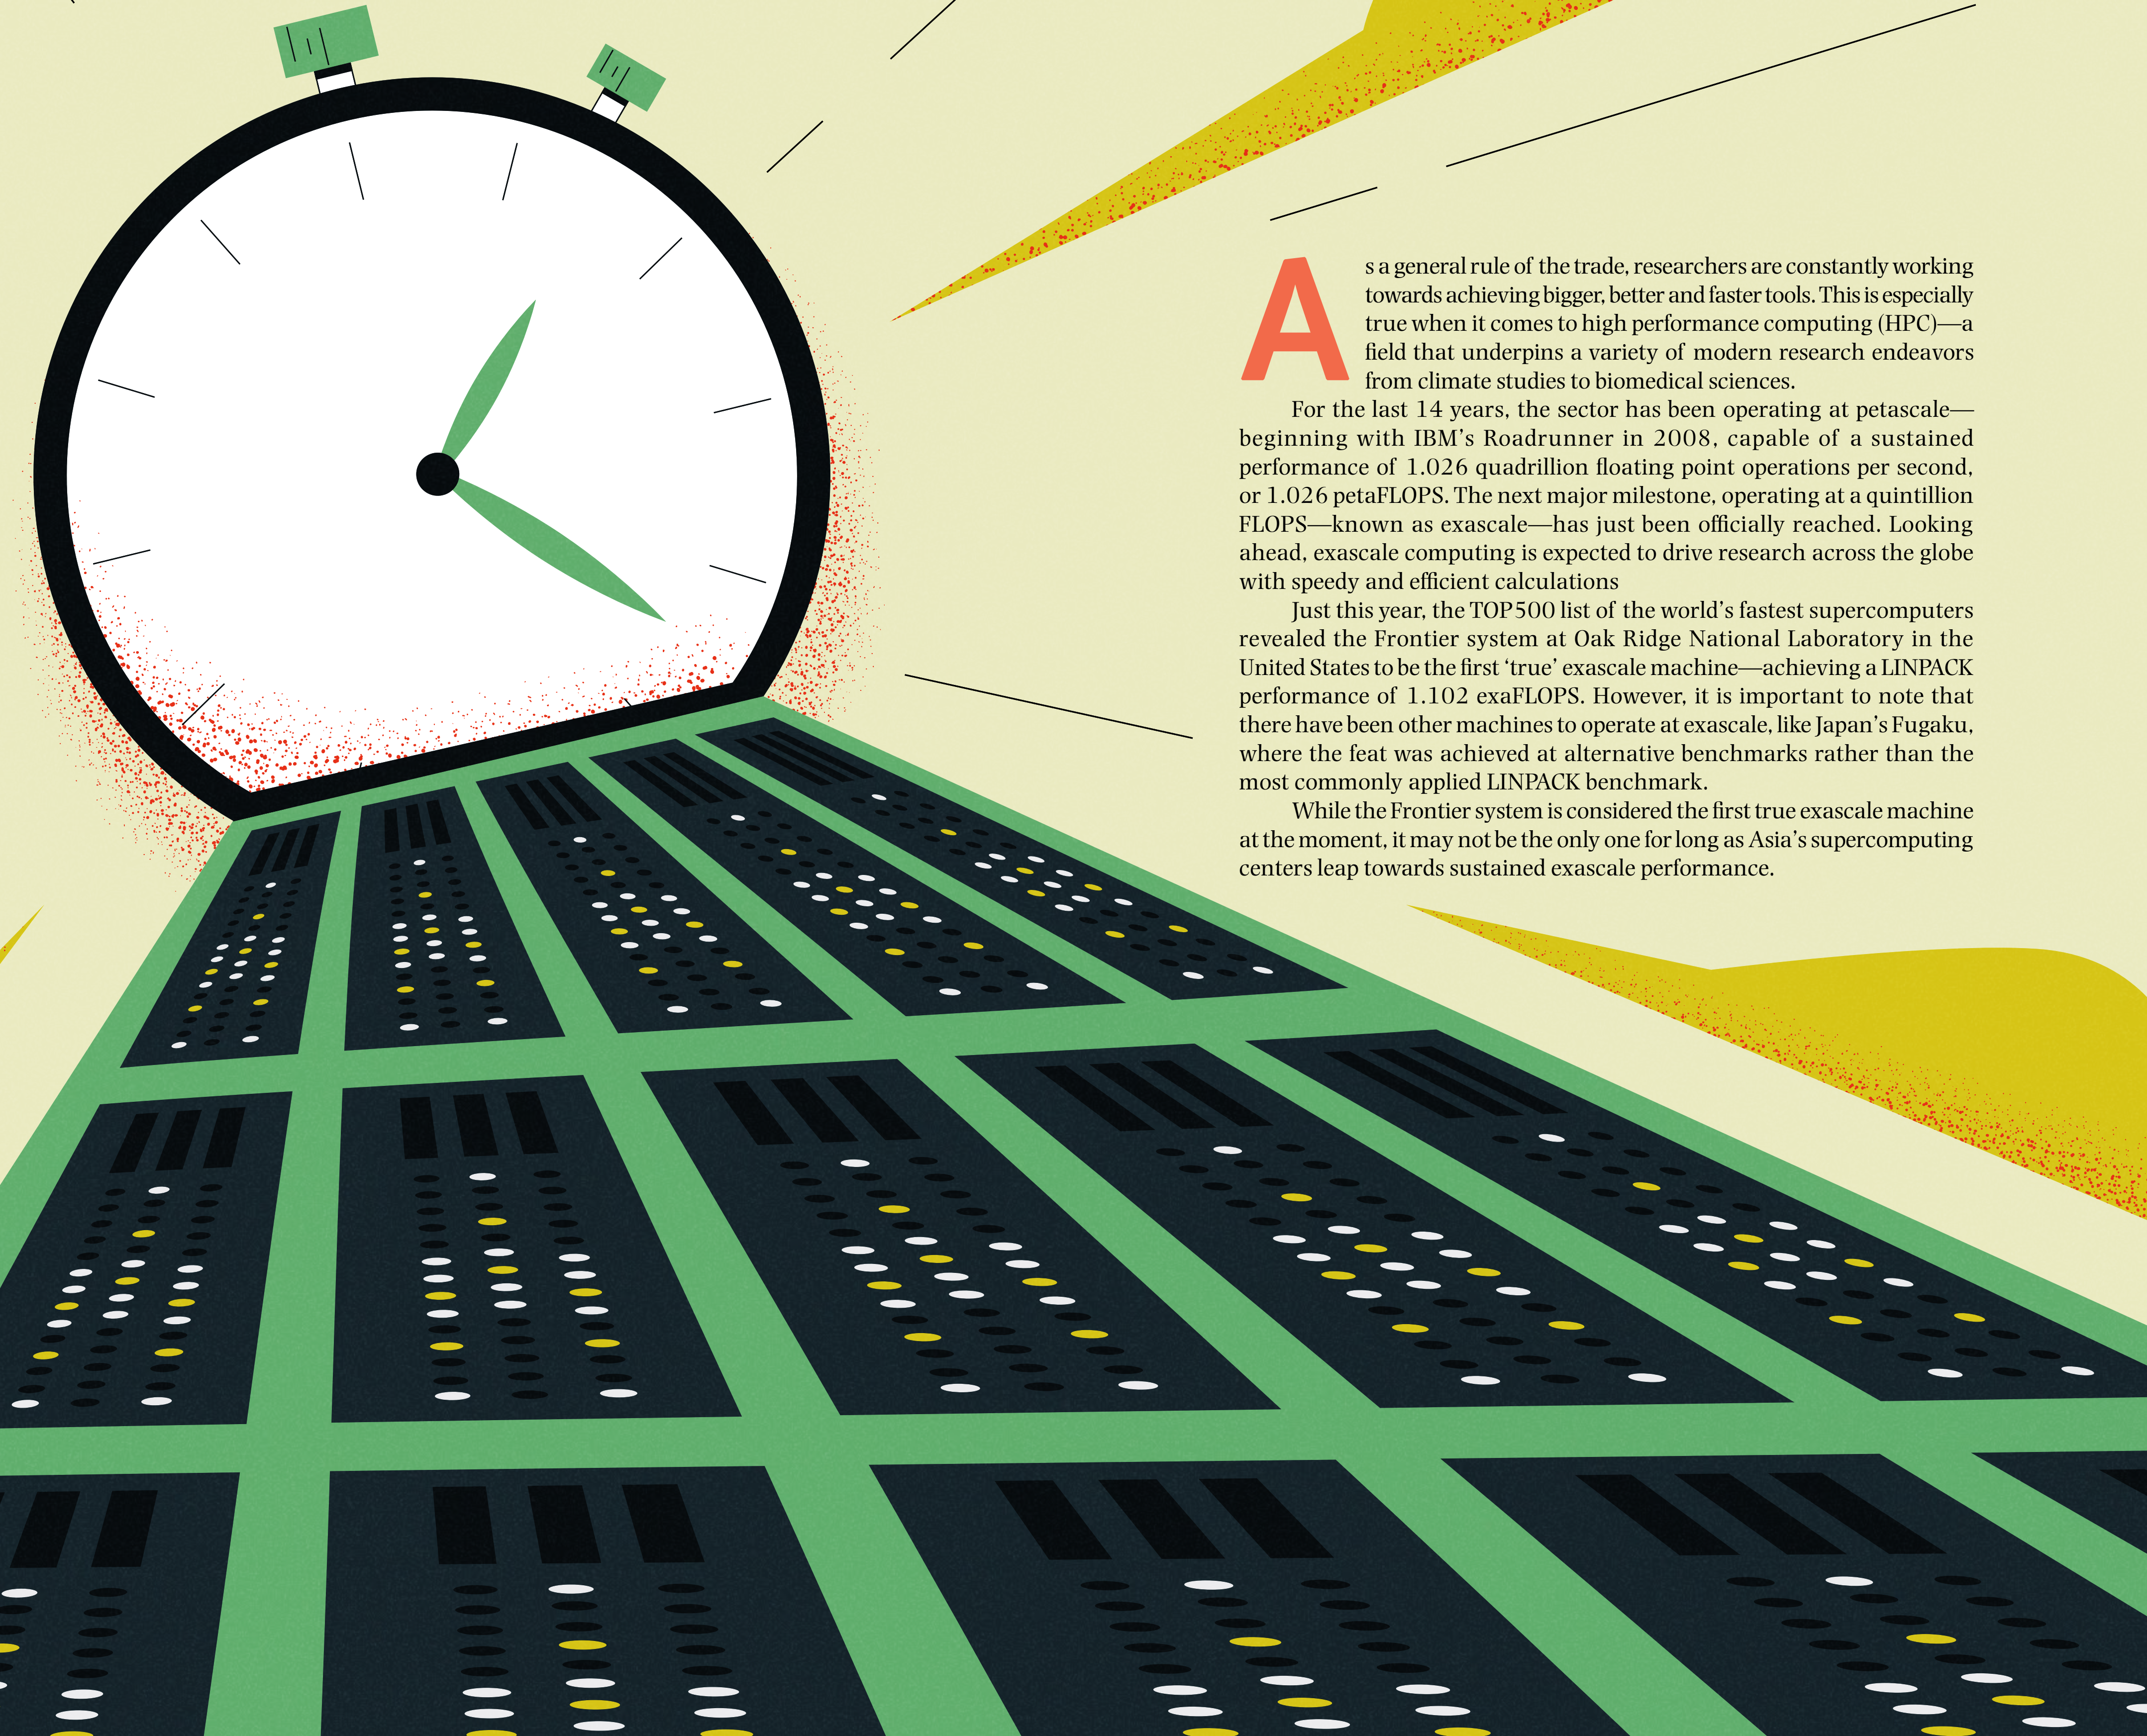

As a general rule of the trade, researchers are constantly working towards achieving bigger, better and faster tools. This is especially true when it comes to high performance computing (HPC)—a field that underpins a variety of modern research endeavors from climate studies to biomedical sciences.

For the last 14 years, the sector has been operating at petascale—beginning with IBM's Roadrunner in 2008, capable of a sustained performance of 1.026 quadrillion floating point operations per second, or 1.026 petaFLOPS. The next major milestone, operating at a quintillion FLOPS—known as exascale—has just been officially reached. Looking ahead, exascale computing is expected to drive research across the globe with speedy and efficient calculations

Just this year, the TOP500 list of the world's fastest supercomputers revealed the Frontier system at Oak Ridge National Laboratory in the United States to be the first 'true' exascale machine—achieving a LINPACK performance of 1.102 exaFLOPS. However, it is important to note that there have been other machines to operate at exascale, like Japan's Fugaku, where the feat was achieved at alternative benchmarks rather than the most commonly applied LINPACK benchmark.

While the Frontier system is considered the first true exascale machine at the moment, it may not be the only one for long as Asia's supercomputing centers leap towards sustained exascale performance.

THE IMPORTANCE OF APPLICATION

Operating at a LINPACK benchmark of 442 petaFLOPS, Japan's Fugaku supercomputer previously held the top spot on the TOP500 list from June 2020 to November 2022. Housed at the RIKEN Center for Computational Science and jointly developed by Fujitsu and RIKEN, the machine has played a significant role in research in Japan—leading to breakthroughs in medicine and climate modeling.

For example, it was announced that with Fugaku, researchers were able to develop a new artificial intelligence (AI) system capable of swiftly predicting tsunami flooding on regular computers. In January 2021, the supercomputer was also used to simulate the movement of molecules and determine how proteins interact with roughly 2,000 existing drugs to find an effective treatment for COVID-19.

In such applications, supercomputers can operate at reduced precision to increase performance while maintaining sufficient accuracy—a state in which Fugaku has achieved performance above one exaFLOP.

"Typically, there are hundreds, even thousands, of jobs running on Fugaku," explained Professor Satoshi Matsuoka, director of the RIKEN Center for Computational Science, in an interview with *Supercomputing Asia*. "It's hard to determine if it's operating at full exascale at any one time, but it's definitely operating at a level that's expected of exascale machines."

Significantly, Matsuoka describes the Fugaku supercomputer as a general-purpose machine—with portions of it available to researchers and commercial users all over the world for a variety of uses. As such, projects usually don't take up very much of the machine and may not demonstrate its full power.

However, while it is clear that meeting the needs of researchers exceeds the symbolic need to reach specific benchmarks—it remains important to occasionally demonstrate the limits of the supercomputer.

"We must not only test the limits of the methodology, but also determine how fast the machine goes at a full scale run with half or more of it in use," said Matsuoka.

Since pushing the envelope with Fugaku's capabilities, the team at RIKEN has begun working on its more powerful successor slated to be launched by end of the decade. "This machine will be a considerable effort involving all the major supercomputing stakeholders in Japan, as well as entities and major companies abroad," shared Matsuoka.

MORE THAN A DROP IN THE OCEAN

Operating at 93 petaFLOPS, China's Sunway Taihu Light at the National Supercomputing Center in Wuxi remains in the top 10 of the world's most powerful supercomputers. However, it seems that China has more in the works. The nation is thought to be operating two exascale supercomputers—neither one has been officially publicly disclosed.

The successor to the Sunway TaihuLight, known as the Sunway Oceanlite is reported to have reached a 1.3 exaFLOPS peak performance after being benchmarked in March 2022.

According to available research, the machine is already in use and plays a starring role in a recent project designed to approach brain-scale AI where the number of parameters is similar to the number of synapses in the human brain. In fact, the project is the first to target training brain-scale models on an entire exascale supercomputer, revealing the full potential of the machine.

So far, it has been reported that the largest tested configuration of the OceanLight system accessed 107,520 notes for an impressive 41.93 million cores across 105 cabinets.

Meanwhile, at the National University of Defense Technology, China holds another supercomputer, Tianhe-3, potentially capable of performing at roughly the same speeds. Similarly deployed for training deep neural networks, Tianhe-3 operates on fully domestic architecture. Specifically, the machine is based on the FeiTing line of processors from Phytium. As China continues to forge ahead as a leading HPC center, it does so independently and relies on native architectures to build bigger and faster machines.

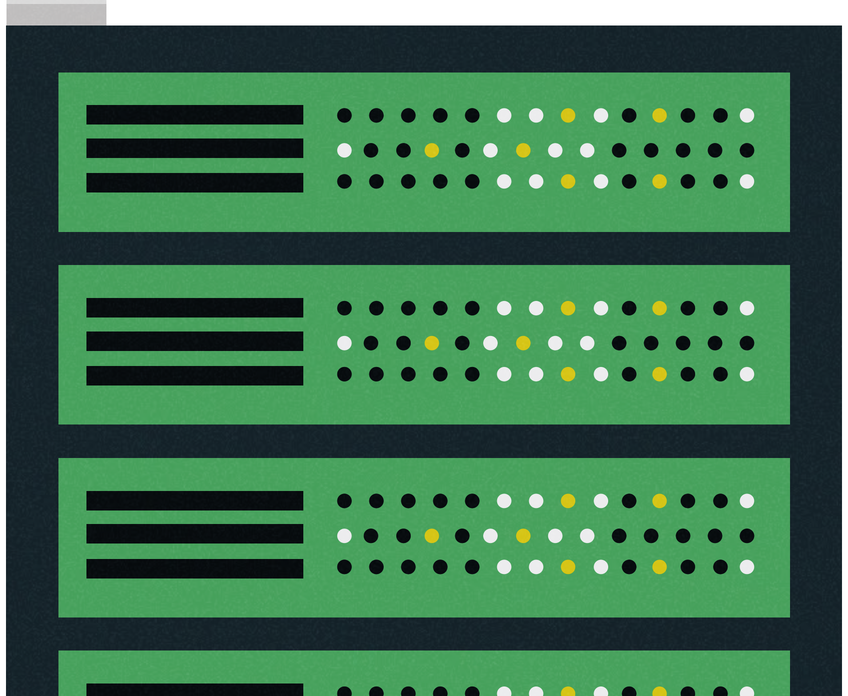

HPC FOR HEALTHCARE

Leaps in supercomputing are also happening in Southeast Asia. At the end of 2021, the National Supercomputing Centre (NSCC) Singapore and National University Health System (NUHS) announced an agreement that would lead to the deployment of supercomputing resources in Singapore's healthcare system. The new machine, named PRESCIENCE, is a petascale supercomputer that will be dedicated towards providing better healthcare and improving the current information infrastructure in Singapore's public healthcare system.

An abundance of patient data is generated and kept by public health systems globally, with an estimated average of 50 petabytes of data produced annually. With this, hospitals are in a prime position to make use of this data using machine learning and AI to provide further insight into a patient's outcome and give the best course of treatment to improve their wellbeing if the patient's health is deteriorating—NUHS plans to do just that with PRESCIENCE.

The current collaborative agreement with NUHS is part of the NSCC's larger roadmap towards growing and maturing Singapore's HPC research landscape. More recently, SingHealth, which is also part of Singapore's healthcare ecosystem, signed agreements with NSCC and Nvidia to deploy a supercomputer to advance medical research within their healthcare cluster.

THE ROAD AHEAD

Earlier this year, India's National Supercomputing Mission (NSM) announced the deployment of "PARAM Ganga", a new petascale supercomputer part of the "PARAM" ("supreme" in Sanskrit) series of machines at the Indian Institute of Technology (IIT) Roorkee. Operating at 1.66 petaFLOPS, this machine is one of a whopping nine new supercomputing systems slated to launch this year.

The NSM, launched in 2015, has deployed 15 systems across the nation since its inception. With many more machines planned, this follows the Indian government's ambitious goal of developing its own supercomputing systems built with parts developed and manufactured in India. Development of indigenous technology is part of the Indian government's bigger goal of furthering supercomputing research and boosting India's own research capabilities across the public and private sectors.

Eventually, upscaling from current petascale computing to exascale computing via the PARAM series is the next step forward for India. This goal is inching ever closer with the planned launch of the PARAM SHANKH exascale system in 2024, which would overtake India's current most powerful machine the PARAM Siddhi-AI, clocking in at a LINPACK performance of 4.62 petaFLOPS.

Between now and the PARAM SHANKH launch, India's researchers and engineers are working on developing and improving native architecture to support exascale computing. International partnerships with Intel Foundry Services and Taiwan Semiconductor Manufacturing Company are further driving India's progress to exascale computing through collaboration.

As Asia continues to make great strides, so too does the rest of the world. Over in the EU, Finland's LUMI system now occupies the third spot in the TOP500 list clocking in a LINPACK performance of 151.9 petaFLOPS. In addition to speed, LUMI and the US' Frontier system are focused on power efficiency as well, with both reaching the top three in the TOP500's GREEN500 list.

However, in testing and ranking the limits of supercomputers worldwide, it is important to remember the reason for their existence—not just to compete for speed, but to aid researchers efficiently and effectively as they collaborate to seek solutions to the world's most pressing challenges.

> "We must not only test the limits of the methodology, but also determine how fast the machine goes at a full scale run with half or more of it in use."
>
> **Professor Satoshi Matsuoka**
> Director of RIKEN's Center for Computational Science

Business Bytes

TAIWAN EYES HPC BOOST FOR HEALTHCARE

Taiwan will soon build its first-ever artificial intelligence (AI) supercomputer for biomedical research, thanks to a new partnership between the National Health Research Institutes (NHRI), Taiwan-based computer manufacturer Asus and US-based chipmaker Nvidia Corporation.

High performance computing (HPC) has a multitude of applications in the medical and pharmaceutical sectors, including genetics and multi-omics analysis, vaccine development and medical imagery identification. AI-powered clinical services are also set to be game-changers in enhancing human health and well-being, ranging from more accurate diagnostics to highly personalized treatment selection.

Taiwan saw the need to bolster its biomedical research with HPC capabilities, tapping Nvidia's supercomputing, AI modeling and analytical computing technologies as well as ASUS' cloud platform. Additionally, ASUS and NHRI will join forces to accelerate the development of AI-powered software focusing on big data for healthcare applications.

By combining these advanced computing capacities with the wealth of clinical data from NHRI, Taiwan's biomedical sector will be well-equipped to push the frontiers for healthcare and innovation in tandem.

ACCELERATING COMPUTING AND MODELING AT CityU

A new HPC cluster built by global edge-to-cloud company Hewlett Packard Enterprise (HPE) has recently arrived on campus at the City University of Hong Kong (CityU). Named CityU Burgundy, this supercomputing cluster can deliver nearly ten times faster computing speed over the university's previous HPC resources to advance research discoveries from biomedicine to behavioral science.

CityU Burgundy features HPE Apollo 2000 and HPE Apollo 6500 Gen10 Plus systems, which are purpose-built and density-optimized platforms designed for demanding HPC and AI applications. The new cluster is also equipped with 328 AMD EPYC™ 7742 processors, 56 Nvidia V100 Tensor Core graphical processing units (GPUs) and eight Nvidia A100 80-gigabyte Tensor Core GPUs. Together, these components can offer more efficient image analysis, modeling and simulation capabilities that are key in AI and machine learning.

Besides the substantial boost in speed, CityU Burgundy is also set to support inclusive and collaborative research. It serves as a centralized facility, consolidating the HPC resources into a single location to enable easier access for various stakeholders while reducing space requirements for data centers.

Biomedical researchers are already using the new HPC cluster to integrate genetic and environmental data from diverse populations. Such endeavors seek to better understand the complexity of chronic diseases and explore novel diagnostic and treatment approaches.

With the improved GPU capabilities and lower latency, CityU Burgundy's resources are also being applied to other disciplines that traditionally were not use cases for HPC—including new data visualization projects, analytics for public policy and consumer behavioral science.

"The new HPC cluster brings us one step closer to building Hong Kong's most powerful HPC platform for academia while achieving operational efficiency and reduced costs," said Dr. Dominic Chien, senior scientific officer (HPC), in a press release. "HPC plays a critical role in helping us build and support a world-class research team to continue making scientific breakthroughs for humankind."

LEVELING UP SINGAPORE'S QUANTUM ECOSYSTEM

As a thriving hub for digital innovation, Singapore has further invested in its quantum sector, launching three national platforms to take the country's technological prowess to new heights. The three initiatives are the National Quantum-Safe Network (NQSN), National Quantum Computing Hub (NQCH) and National Quantum Fabless Foundry (NQFF).

Under Singapore's Research, Innovation and Enterprise 2020 (RIE2020) plan, the Quantum Engineering Programme of the National Research Foundation Singapore (NRF) has pledged at least S$23.5 million to support these platforms for up to three and a half years.

Onboard the NQSN project are over 15 private and government collaborators, including the Infocomm Media Development Authority (IMDA), with the initiative led by the Centre for Quantum Technologies (CQT) teams at the National University of Singapore, as well as Nanyang Technological University, Singapore. The NQSN will conduct nationwide trials on quantum-safe communication technologies, aiming to achieve robust network security to safeguard sensitive data and critical infrastructure.

Together with A*STAR's Institute of High Performance Computing (IHPC), CQT teams are also involved in the NQCH initiative to develop quantum computing hardware and middleware. In support, the National Supercomputing Centre (NSCC) Singapore will host a quantum computing facility and provide resources to explore industrial applications in finance, logistics, chemistry and more.

Meanwhile, the NQFF at A*STAR's Institute of Materials Research and Engineering (IMRE) will develop microfabrication techniques to manufacture component materials and quantum devices for computation, communication and sensing applications.

"The launch of the three national platforms signals the intent and ambition of Singapore to build upon our past investments in quantum technologies, and take it further through close industry development with our partner agencies," shared NRF CEO Professor Low Teck Seng, to the press.

By fostering talent development and forging public-private partnerships, these national quantum initiatives are envisioned to create an enabling environment for quantum innovation across Singapore.

SCALING THE SUPERCOMPUTING SECTOR

Steady growth is on the horizon for the HPC industry in the coming decade, after global revenue hit US$42 billion worldwide in 2021. Market research and consultancy firm Emergen Research has forecasted a 6.2-percent compound annual growth rate (CAGR) to push the market size to US$71.8 billion by 2030.

Meanwhile, according to technology research and advisory company Technavio, the Asia-Pacific region is expected to account for 49 percent of the global HPC market's growth from 2021 to 2026. Technavio also predicted more significant progress, pegging the CAGR at 11.31 percent and an overall revenue increase of US$27.15 billion over the forecast period.

In particular, the cloud computing and data center segments are likely to become key drivers of the HPC sector's growth. Cloud computing can potentially encourage HPC adoption among small and medium-sized enterprises by connecting them to otherwise high-cost and inaccessible supercomputing resources.

With big data analytics on the rise, supercomputer servers and data centers are needed to store and process voluminous amounts of data—ultimately to generate meaningful insights and accelerate digital transformation across industries. As such, this segment is also predicted to gain further ground among enterprises and governments, spurring substantial HPC market growth over the decade.

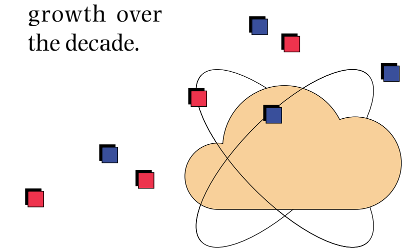

Super Snapshot

ASPIRE 2A IN NUMBERS

Singapore's newest and fastest national supercomputer will supercharge scientific research across Singapore

The Advanced Supercomputer for Petascale Innovation Research and Enterprise (ASPIRE 2A) is the first in a series of supercomputers that will be deployed in phases to expand and upgrade Singapore's high performance computing (HPC) capabilities. When operational, ASPIRE 2A is expected to provide an aggregate of up to 10 petaflops of raw compute power and be eight times more powerful than the current ASPIRE 1 supercomputer. ASPIRE 2A's computing power will be exploited for the advancement of climate and biomedical research, as well as smart nation initiatives by the government, research institutions and companies in Singapore.

How ASPIRE 2A stacks up against ASPIRE 1

> 3.5x more Cores

5x more Compact (1.5x less Nodes)

2x more GPUs

8x Compute Power

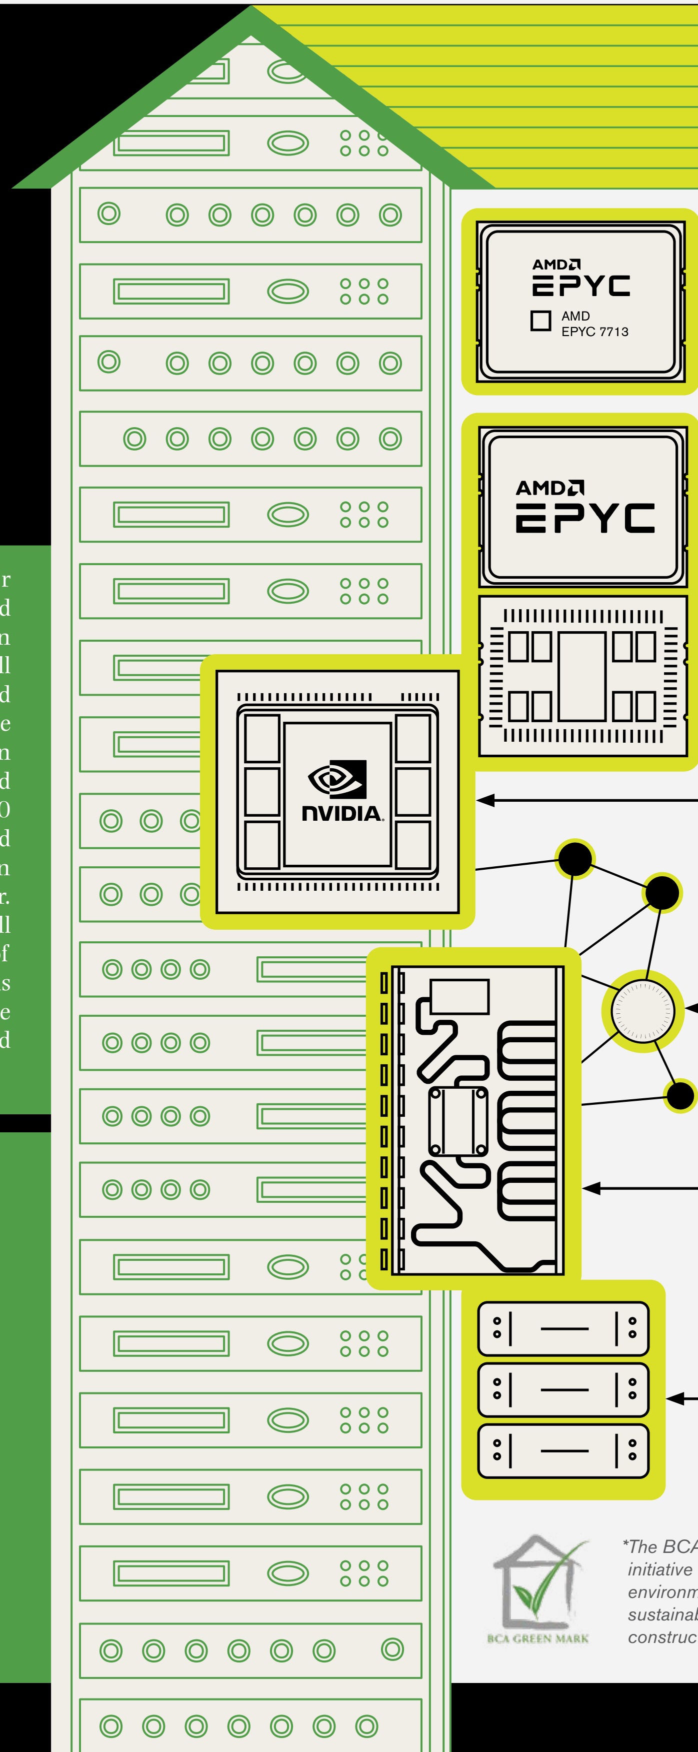

105,984 Cores
CPU (AMD EPYC™ 7713)
800 Nodes

1,024 Cores
High Frequency Nodes (AMD EPYC™ 75F3)
16 Nodes

352 GPUs
Accelerated Nodes GPU (Nvidia A100)
82 Nodes

476 TB
Total System Memory

25 PBytes
Storage (Spinning + Nearline)

10 PBytes
Scratch Disk

BCA GREEN MARK

**The BCA Green Mark Scheme is an initiative to create a more sustainable built environment in Singapore by promoting sustainable design, and best practices in construction and operations in buildings.*

Smarter champions environmental change.

Author
Nigel Lee
Country General Manager of Lenovo Singapore

Drive sustainability with green products and solutions engineered to reduce emissions.

A recent study by Engie Impact, a sustainability solutions company, found that 90% of Asia-Pacific companies lack confidence in their decarbonization capabilities. And KPMG's 2021 Net Zero Readiness Index ranks Singapore at the 15th spot, out of 25 countries in terms of zero emission goals. These are alarming figures, given how poor Environmental Social Governance can weaken a company's reputation, and incur a high carbon tax. The good news is that enterprises are implementing strong sustainability measures. Yet, few can do it on their own.

With sustainability ingrained in everything we offer, at Lenovo, we go the extra mile to support our customers in their sustainability journey. It should come as no surprise, therefore, that we are one of the first PC manufacturers to offer CO_2 Offset Services to our commercial customers. With the aim of halving carbon emissions by 2030, this service helps customers meet their carbon reduction objectives, reflecting Lenovo's commitment to Environmental Social Governance. It tracks and controls emissions at every stage of the product life cycle, from the factory floor to the end destination. In fact, we are pleased to note that, to date, our business customers alone have offset over 50,000 metric tonnes of CO_2.

Our focus on the environment even extends to our products, with the new Z Series ThinkPads - The Z13 and Z16, powered by AMD Ryzen™ PRO technology. Manufactured from recycled materials, they reflect our growing commitment to driving change.

ThinkPad Z13 & Z16 laptops

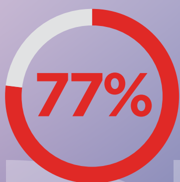

of consumers say it's important that brands are sustainable and environmentally responsible[1]

To learn more, please contact
Alexis.Lee@amd.com

[1] Climate change and corporates: Past the tipping point with customers and stockmarkets. Deutsche Bank Research, 2019