AsianScientist
Issue 13
January 2023

SUPERCOMPUTING ASIA

CAPPING CARBON ON ASIA'S SUPERCOMPUTERS

WITH GREAT COMPUTING POWER COMES GREAT ENERGY LIABILITY—UNTIL NOW

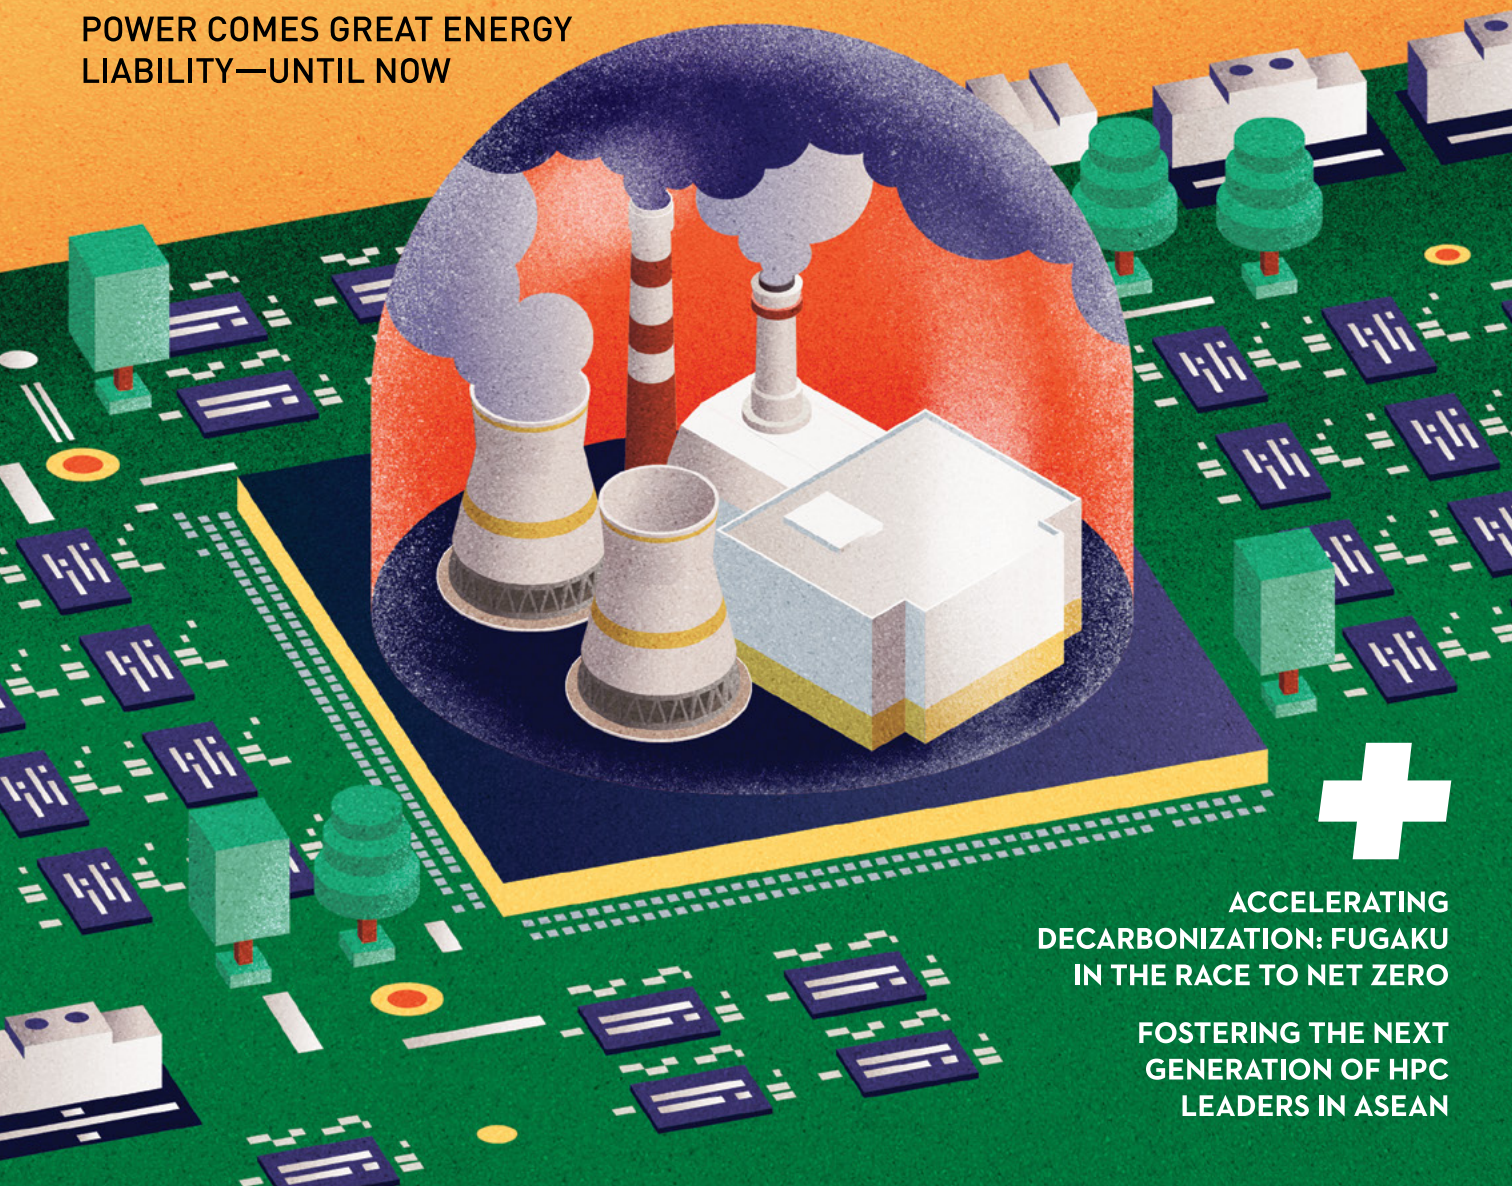

ACCELERATING DECARBONIZATION: FUGAKU IN THE RACE TO NET ZERO

FOSTERING THE NEXT GENERATION OF HPC LEADERS IN ASEAN

CONTENTS

Issue 13
January 2023

FEATURES

p. 26

PULLING THE CURTAINS ON CLIMATE CHANGE

Understanding our changing climate

COVER STORY

p. 18

Capping Carbon On Asia's Supercomputers

With great computing power comes great energy liability—until now

FEATURES

p. 12

Accelerating Decarbonization: Fugaku In The Race To Net Zero

The centerpiece of Japan's Society 5.0

p. 32

Fostering The Next Generation Of HPC Leaders In ASEAN

Creating valuable connections

p. 06

Digital Dispatch

Supercomputing news from around the world

p. 36

Business Bytes

The latest industry moves

p. 38

Super Snapshot

Going farther, together: Collaborating for a quantum future

EDITOR'S NOTE

Decarbonization is more than just a buzzword—the reduction of carbon emissions is one of the key solutions to climate change, the biggest existential crisis we face today, and supercomputers are helping us achieve it.

In this January 2023 issue of *Supercomputing Asia*, we take a look at how researchers in the region are using high performance computing to better understand our changing climate (*Pulling the Curtains on Climate Change*, p. 26) as well as reduce and capture carbon emissions (*Accelerating Decarbonization: Fugaku in the Race to Net Zero*, p. 12).

With great computational power comes a great electric bill, and our cover story (*Capping Carbon on Asia's Supercomputers*, p. 18) explores ongoing solutions to the worrying carbon problem that has long accompanied supercomputing ability. These solutions include better hardware and software technology, but also more purposeful governance and policies that can make supercomputers no longer a contributor to carbon emissions but rather, a key part of the solution.

Lastly, we take a glimpse into the future of supercomputing as we explore events like the EU-ASEAN HPC School (*Fostering the Next Generation of HPC Leaders in ASEAN*, p. 32), a week-long course on high performance computing for young researchers.

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

ASSISTANT EDITOR
Marinel Mamac

CONTRIBUTORS
Jihan Al-Shdifat
Tristan Mañalac
Mitchell Lim

ART DIRECTOR
Shelly Liew

SENIOR DESIGNER
Ajun Chuah

DESIGNERS
Yipei Lieu
Jorgina Tan
Wong Wey Wen

SALES & MARKETING
Samantha Yeap
Kata Llamas
Audrey Tan

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

Digital Dispatch

HIGH PERFORMANCE COMPUTING SHEDS LIGHT ON LUNAR MYSTERY

Most scientists agree that our moon was created 4.5 billion years ago, when a large body the size of Mars collided with Earth. In this theory, the body—dubbed Theia—was destroyed, its pieces becoming debris in our Earth's orbit. The moon today is said to be made of those broken pieces, which coalesced over time.

To test out this theory, astronauts from the US Apollo program collected lunar rock samples that had structural and isotopic similarities to early Earth rock samples.

To better understand our moon's origin, researchers from the UK's Durham University, armed with the power of the DiRAC supercomputing facility, ran a range of high-resolution simulations that incorporated key parameters like the angles, speeds and masses of Earth and Theia.

In a study published by *The Astrophysical Journal Letters*, they found that Theia's moon-like mass coalesced immediately in our orbit following its collision with Earth, and brought with it proto-Earth material that was ejected from our planet in the impact. The entire process took hours, not days, and is why there are similarities between the rock samples collected from the moon and Earth.

Astronauts from NASA's Artemis Program are set to collect fresh lunar samples to test this new theory in the coming decades.

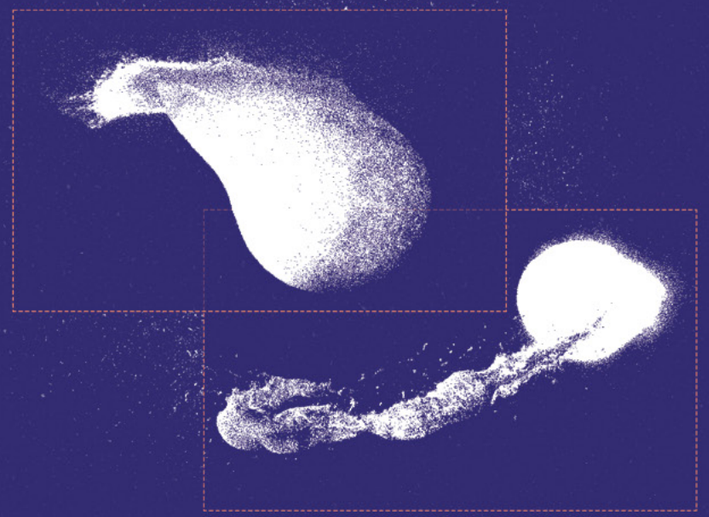

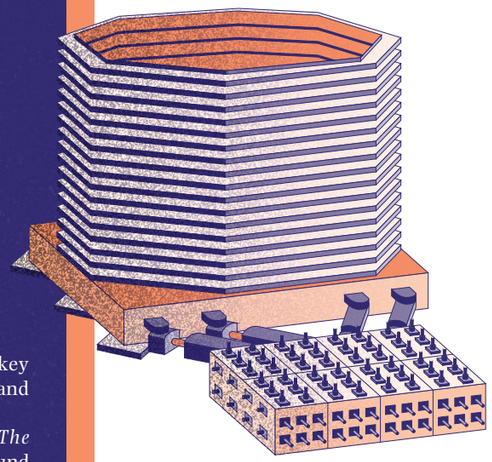

SUPERCHARGING EARTHQUAKE MODELING

Some of the world's top supercomputers—namely, Berkeley Lab's Perlmutter as well as Oak Ridge National Laboratory's Summit and soon, Frontier, the world's first official exascale supercomputer—are crunching the numbers to help researchers better understand earthquakes and develop more earthquake-resilient infrastructure.

Scientists are deploying high performance computing in two related projects. The first is the Large-Scale Laminar Soil Box System, which uses a hydraulic platform capable of exerting 1.25 million pounds of force to shake up to 350 tons of soil and simulate the effect of earthquakes. The second project is the ongoing EQSIM program, which simulates earthquakes on supercomputers while also learning from the data gathered by the Soil Box System.

The resulting data sets will be made available to broader research communities on the open-access simulation database of the Pacific Earthquake Engineering Research (PEER) Center.

CONGRATULATIONS TO THE WINNERS OF SUPERCON2022!

Organized by Tokyo Tech's Global Scientific Information and Computing Center (GSIC), Osaka University's Cybermedia Center (CMC), and the RIKEN Center for Computational Science (R-CCS), the 2022 Supercomputing Contest (SuperCon2022) for high school and technical college students recently announced the following winners:

1ST PLACE
Team Citrus
Nada Senior High School

2ND PLACE
Team noCrepuS
Shibuya Kyoiku Gakuen Makuhari Senior High School

3RD PLACE
Team bubuzuke
Matsumoto Shuho Secondary Education School

Team T0M0K4ZU
Tokyo Tech High School of Science and Technology

For winning the contest, Team Citrus also received the Minister of Education, Culture, Sports, Science and Technology Award as well as the Academic Award from the Institute of Electronics, Information and Communication Engineers, and Information Processing Society of Japan. This is the team's second consecutive year winning the SuperCon.

Now on its 28th edition, the SuperCon2022 finals challenged its participants to "distinguish strings with an automaton." The main problem was inspired by a famous unsolved problem in automata theory, a branch of theoretical computer science known for its various applications in pattern matching and programming languages.

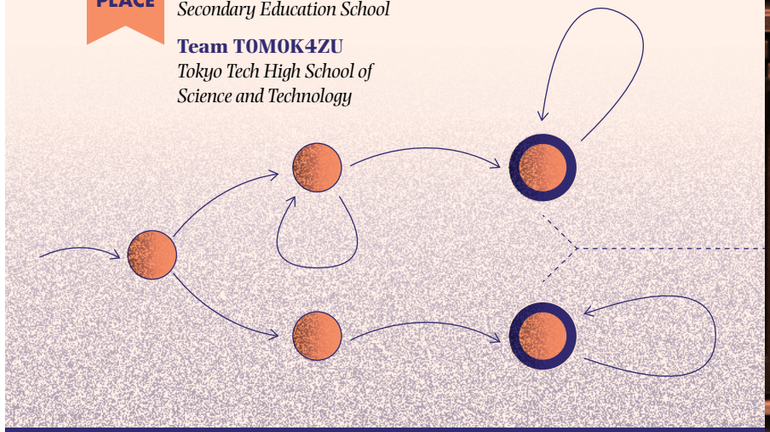

Digital Dispatch

AUSTRALIA INTRODUCES TAINGIWILTA, THE POWERFUL

In August 2022, the Australian Department of Defense unveiled a new supercomputer called Taingiwilta, whose name means "powerful" in the language of the Kaurna people of South Australia. The new supercomputer will be used to design, develop and analyze national security systems as well as modern weapon systems, including nuclear-powered submarines.

Taingiwilta's work and specifications are classified, though the Australian government has said that it is among the top 50 most powerful supercomputers in the world.

What is available to the public, however, is the artwork on Taingiwilta, which was created by three generations of a Kaurna family. The artwork is called Ngangk Mukarta Warpu Yerta, which translates to "Women's intelligence on Earth."

WHAT'S UP!

SUPERCOMPUTINGASIA 2023

SupercomputingAsia 2023 (SCA23) is co-organized by different supercomputing centers in the region, such as those in Australia, Japan, Singapore and Thailand, and anchored by the National Supercomputing Centre (NSCC) Singapore. It aims to promote a vibrant and shared HPC ecosystem for both the public and private sectors in Asia.

For the first time ever, HPC Asia 2023 will be co-located and held in conjunction with SCA23 in Singapore. HPC Asia, or the International Conference on High Performance Computing in the Asia-Pacific Region, is an international conference series for the Asia Pacific region on HPC technologies, which aims to foster an exchange of ideas, research results and case studies related to all issues of HPC.

In addition, the Conference on Next Generation Arithmetic (CoNGA), the leading conference on emerging technologies for computer arithmetic, will also be back. It will be held in conjunction with SCA23. Gain the latest news and updates on the developments of breakthroughs with next generation data formats and their corresponding hardware, tools, applications and services and exchange ideas on what next generation arithmetic should be.

This year's theme is "Sustainable Supercomputing for a Greener Future."

Book your ticket to SCA23 to be part of this historic event. Discover the latest HPC trends, delivered straight from thought leaders in academia and industry!

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

WHERE

SINGAPORE

WHEN

FEBRUARY 27 – MARCH 2, 2023

ISC HIGH PERFORMANCE 2023

The International Supercomputing Conference (ISC) High Performance 2023 is back this year with a theme of "Imagine Tomorrow," at the Congress Center Hamburg in Germany from May 21 to 25. Like last year's iteration, the ISC High Performance 2023 conference and exhibit will be an on-site event with comprehensive on-demand content for individuals who cannot attend in person.

The event, the 38th of its kind, aims to be an active collaborative exchange on five key areas in HPC technology: system architecture; parallel programming models and performance modeling; applications and algorithms; machine learning; and quantum computing. As Moore's Law gradually slows, the ISC High Performance 2023 program looks forward to a new era for high performance computing—looking beyond exascale and towards emerging technologies and economic models.

First held in 1986, ISC High Performance distinguishes itself as the world's oldest and Europe's most significant forum for the HPC, machine learning and high performance data analytics communities.

Save the date and get ready to join the 3,000-strong community of attendees, speakers and exhibitors in this year's conference!

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

WHERE

HAMBURG, GERMANY

WHEN

MAY 21 – 25, 2023

Digital Dispatch

SUPERCOMPUTERS FOR THE EARTH:

HPC INNOVATION CHALLENGE WINNERS ANNOUNCED

Organized by the National Supercomputing Centre (NSCC) Singapore with the support of GeoWorks, SGTech and SGInnovate, the Inaugural HPC Innovation Challenge (HPCIC) for the Environment has announced its pioneer batch of winners.

24
teams joined the challenge

2
categories: student and open

10
teams made it to the Development Phase and presented their solutions on Demo Day

7
weeks to develop their solutions with a mentor and ASPIRE 2A, one of the most powerful supercomputers in Southeast Asia

5,000 SGD
maximum cash prize for the winners

6
winners of the inaugural HPCIC

Student category

Winner

Viridis

Nanyang Technological University (NTU) and Singapore Management University (SMU)

The team developed a numerical assessment tool that automatically assigns walkability, bike-ability and drivability scores to evaluate transport flow.

Q What was the highlight of the HPC Innovation Challenge experience for you?

The highlight was Demo Day. We worked very hard over the summer to develop our solution, which we felt had the potential to impact urban planning and the way Singaporeans commute. Presenting our ideas to industry experts and gaining their recognition for our work made it all worthwhile. Of course, winning the challenge was like icing on the cake for the team!

1st runner up

Bellman Ford

Singapore University of Technology and Design (SUTD) & National University of Singapore (NUS)

The team developed an analytics platform that will aid the decision-making process of urban planners in determining which walking paths to improve on and the best way to do so.

Q What was the highlight of the HPC Innovation Challenge experience for you?

The highlight was being able to see all the interesting innovations that both students and companies were working on. We also had the opportunity to interact with the ASPIRE 2A supercomputer through the Digital Sandbox, which was a novel experience for us.

2nd runner up

FRG Team

Singapore University of Technology and Design (SUTD)

The team proposed to use HPC to efficiently design new fluorophores by conducting virtual screening and systematic data analysis. The new method will reduce chemical waste by eliminating the need to carry out trial-and-error methods.

Q Do you have any advice for students who might be interested in entering challenges like the HPCIC, or want to get involved in HPC-powered research?

Aside from our own field of study, HPCs can be utilized in so many other fields. We hope that this competition will encourage more students at different levels to apply HPC methods in their projects.

Open category

Winner

Aleph Technologies

Aleph Digital Technologies Pte. Ltd.

The startup sought to utilize HPC to improve the AI-engine capability and speed of their platform, by suggesting optimal operating parameters. Aleph's technology helps factories reduce energy consumption and CO_2 generation.

Q What was the inspiration for your winning project?

We believe that digital solutions are going to play a key role in our planet's decarbonization journey. This transitional journey is going to start with emission reduction, followed by waste energy recycling and lastly, replacement with renewables and green options. Since a significant amount of global emissions originate in industrial and power sectors, we are on a mission to reduce their energy consumption through optimization.

1st runner up

GeoPulse

HSC Pipeline Engineering Pte. Ltd.

The engineering company sought to utilize HPC to carry out automated subsurface utility mapping by compiling and analyzing large volumes of GPR, drawings and digital twin data.

Q Do you have any advice for researchers who might be interested in entering challenges like the HPCIC, or want to get involved in HPC-powered research?

A civil engineering contractor genuinely has no business taking part in a high-powered computing challenge. While we'll remain an oddity for some time to come, the fact is that there was a spark of an idea, and NSCC's support in de-mystifying supercomputers enabled us to break free from conventional limitations and pursue it to a surprising degree of success. Imagine what people versed in this field can actually do!

2nd runner up

DBF Minds

Digital Blue Foam Pte. Ltd, Singapore

The AI-architecture startup collaborated with Karamba3D to explore ways of generating synthetic datasets towards reducing the carbon impact of new and existing building designs.

Q Do you have any advice for researchers who might be interested in entering challenges like the HPCIC, or want to get involved in HPC-powered research?

Challenges like HPCIC provide a great opportunity for us to use HPC for heavy computational tasks addressing real-world problems. We would recommend the researchers be as ambitious as possible in defining their goals. Just try to imagine the biggest problems in your field and how the huge computational capabilities of HPC can help you to solve them.

ACCELERATING DECARBONIZATION: FUGAKU IN THE RACE TO NET ZERO

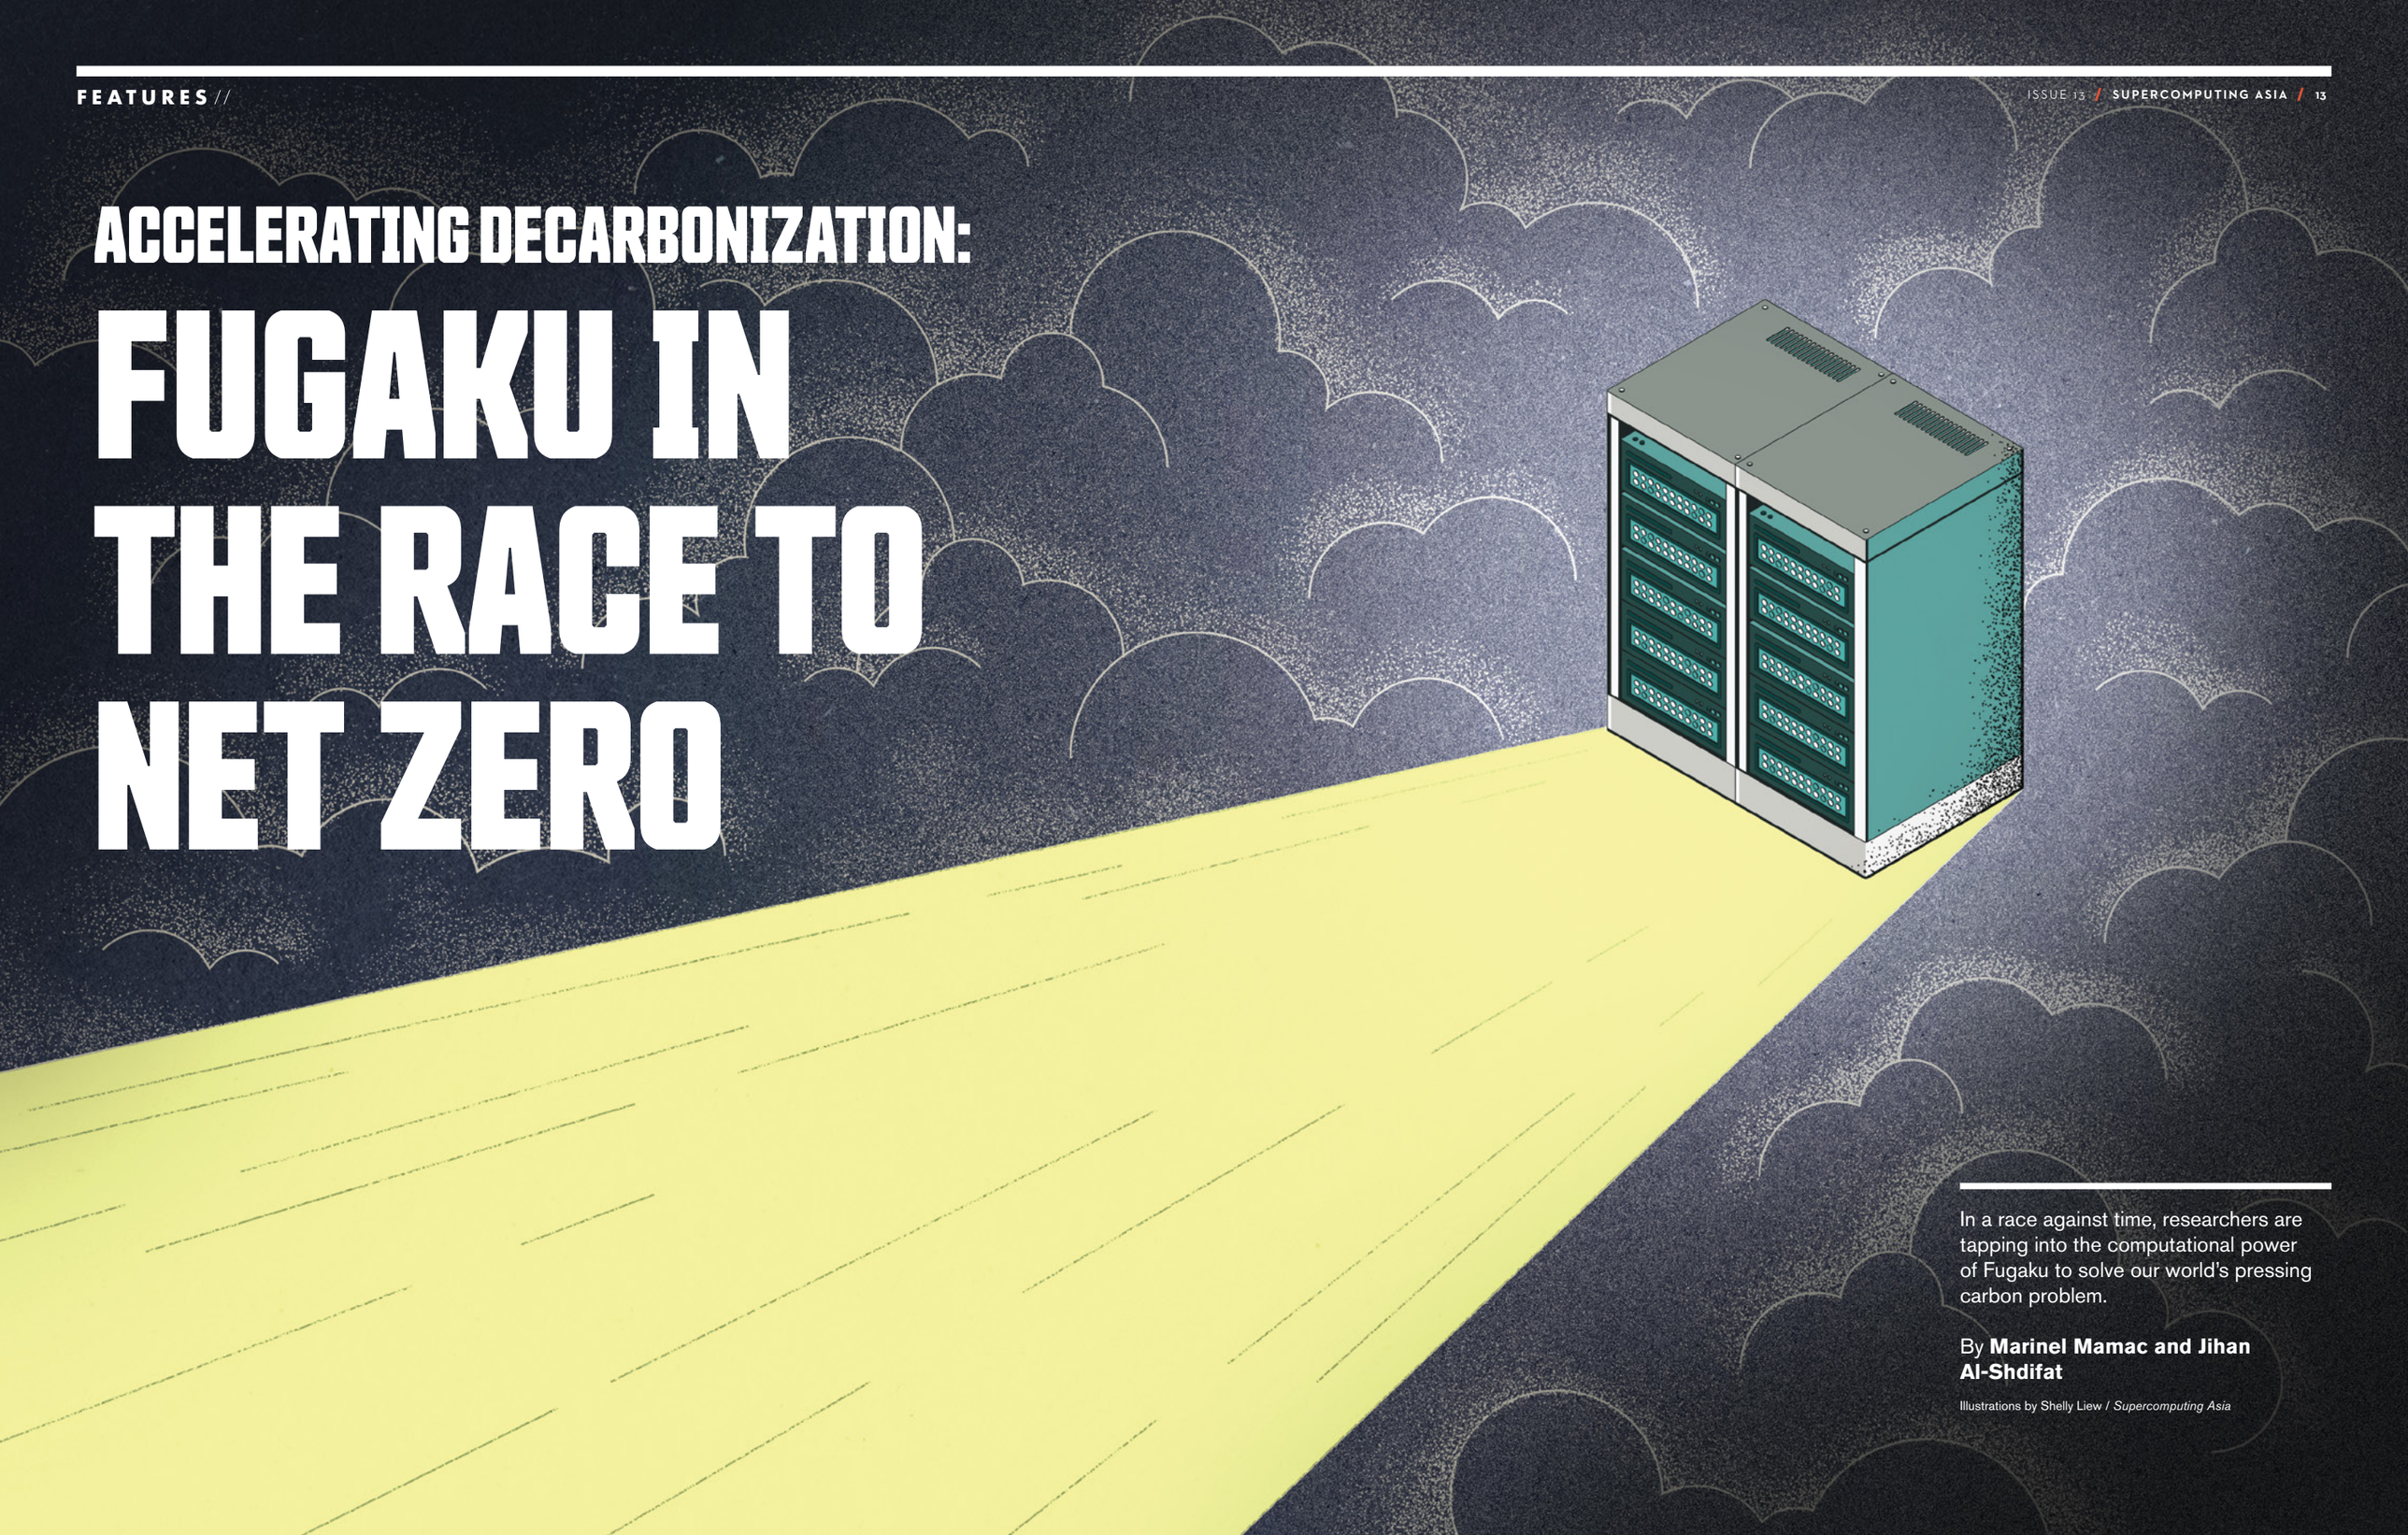

In a race against time, researchers are tapping into the computational power of Fugaku to solve our world's pressing carbon problem.

By **Marinel Mamac and Jihan Al-Shdifat**

Illustrations by Shelly Liew / *Supercomputing Asia*

When it debuted in 2020, the supercomputer Fugaku was hailed as the centerpiece of Japan's Society 5.0—a vision of a country able to solve social problems and advance its economy with the help of digital technology. It was, after all, the world's fastest supercomputer at the time, a title it held on the biannual TOP500 list from its launch until June 2022.

"The main mission of Fugaku is to attain sustainability goals in the areas where it is involved," said Professor Satoshi Matsuoka, director of the RIKEN Center for Computational Science and part of the team behind Fugaku, in an interview with *Supercomputing Asia*.

Achieving 442 petaFLOPS of computational power, Japan's fastest supercomputer was developed by scientists at RIKEN with an application-first philosophy. That is, it wasn't just about achieving computational excellence for its own sake—the machine was built to solve the biggest crises of our time. "Many of these crises pertain to carbon neutrality," Matsuoka noted.

More than just a buzzword, decarbonization has become a crucial goal among researchers in Japan—a country that is among the world's top carbon emitters. By the year 2030, Japan aims to reduce carbon emissions by 46 percent compared to 2013 baseline levels. With today's generation of high performance computing (HPC) hardware, software and talent, Japan is leading Asian countries in achieving this goal.

TACKLING THE PROBLEM OF CARBON

In the half-hour or so that you spend sifting through this magazine, approximately 1.9 million metric tons of carbon dioxide (CO_2) were emitted globally as a result of burning fossil fuel. Of the 50 billion tons of CO_2 and CO_2-equivalent greenhouse gases emitted each year, 73.2 percent comes from the energy sector, 18.4 percent from agriculture and 5.2 percent from industry. The remaining 3.2 percent comes from landfills and wastewater.

Of course, the earth has its own ways of ensuring a balance: Up to 83,000 metric tons of carbon per square kilometer will be sequestered by seagrass meadows, while forests take up around 30,000 metric tons of carbon per square kilometer.

But still, some 40 percent of emissions will make their way to the atmosphere, while 30 percent will be absorbed by seawater, causing ocean acidification. And as carbon emissions continue to rise and we lose our forests to rapid urbanization, wildfires, mining, unsustainable agriculture and rising sea levels, decarbonization innovation has become increasingly crucial.

Tackling the world's carbon problem entails two sides of the same coin. The first involves reducing greenhouse gas emissions—from using renewable energy sources to controlling carbon emissions in agriculture. This also includes ensuring better energy efficiency across different industries. The second involves finding ways to improve the absorption of carbon from the atmosphere, either by capturing emissions directly or enhancing the natural carbon storage of our forests and seagrass meadows.

It's a tall order—but it's one that scientists at RIKEN and the rest of the world are taking on with the power of HPC.

"Each blade is like a skyscraper. These things are so large that the blades sit above the clouds and may actually affect the weather. This makes them very difficult to design."

Professor Satoshi Matsuoka
Director of the RIKEN Center for Computational Science

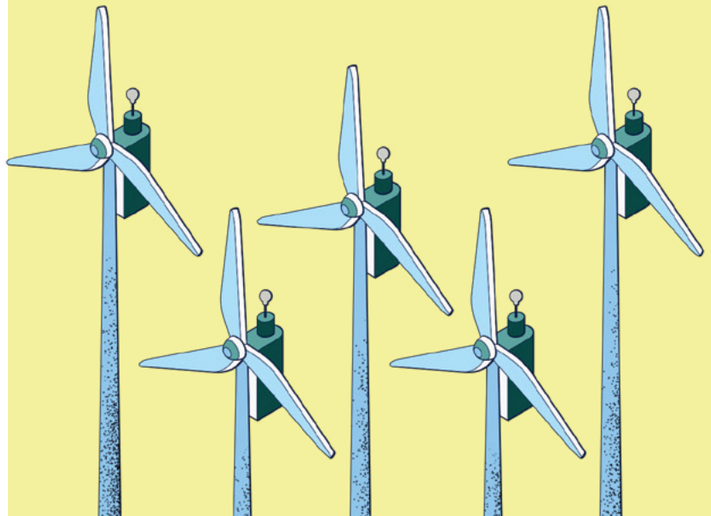

REDUCING CARBON EMISSIONS

Primary among Fugaku's decarbonization research goals is developing a strong pipeline of renewable energy. Matsuoka explained that Japan's biggest hope for carbon neutrality lies in wind and solar energy.

In line with this, the country aims to generate 10 gigawatts of offshore wind power by 2030. To make this possible, the country has been investing in humongous offshore wind farms, with propeller blades as high as 200 meters.

"Each blade is like a skyscraper," shared Matsuoka. "These things are so large that the blades sit above the clouds and may actually affect the weather. This makes them very difficult to design."

To make these offshore wind farms work, scientists have to consider more than just the physics of converting wind energy into electricity. They are using Fugaku's computing power to run simulations based on extreme weather conditions, blade materials and overall design.

As for solar energy, Matsuoka said that Japanese material science teams have been studying how to make safer and more efficient photovoltaic cells locally and, in the process, would need to investigate some 20 million possible substances. "That number is just too huge to run a [physical] experiment," he pointed out.

With Fugaku and artificial intelligence systems designed for solar energy research, scientists can run billions of simulations on all these possible substances, filtering through these compounds much more efficiently. "One of the substances discovered by our team has an energy conversion rate of up to almost 25 percent. This is very promising," Matsuoka said. From there, the next step is to translate these photovoltaic cell simulations to real life and then conduct further research, he said.

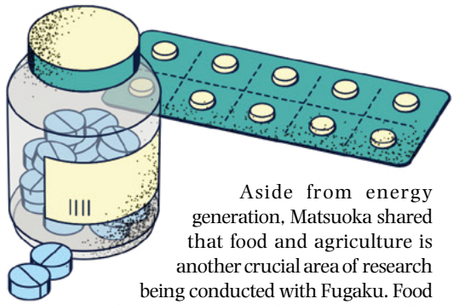

Aside from energy generation, Matsuoka shared that food and agriculture is another crucial area of research being conducted with Fugaku. Food production generates lots of carbon, he pointed out, with cows being the largest producers of greenhouse gases in the sector. Today, Japanese scientists are trying to develop a breed of cows that produce less methane, while others are investigating better and more efficient ways to produce food.

"What's interesting for me is that the technologies these teams are using are very similar to technologies we have already developed for other areas of biology, like human medicine," shared Matsuoka. "Some of the infrastructure we have built to accelerate drug discovery can be translated into food production, because underneath they're using various genomic and proteomic technologies."

Another area that Japanese scientists are looking into is maritime transport. Today's ships are the backbone of global trade, but they also consume a significant amount of the world's total energy—prompting scientists at RIKEN to study how we could make more efficient ships.

Matsuoka explains that designing these ships has always been a huge challenge. In conventional development, maritime design requires scale models. Huge pools are built to test small ship prototypes, with some pools being hundreds of meters in length.

"For the first time, we have the capacity to assist scientists in the maritime design process, eliminating the need for these huge pools," Matsuoka shared. Scientists at the University of Tokyo have begun using Fugaku to conduct pool simulations, taking into account thousands of moving factors like the viscosity of water, the design of the cargo ship and the screws that will propel it. This allows them to boost ships' energy efficiency by 10–15 percent, which Matsuoka pointed out can make a huge impact globally.

"We haven't done a cumulative assessment of just how much potential Fugaku has allowed us to tap into. Maybe that's something we should do, collectively, so we can continue to significantly reduce Japan's overall carbon footprint."

Professor Satoshi Matsuoka
Director of the RIKEN Center for Computational Science

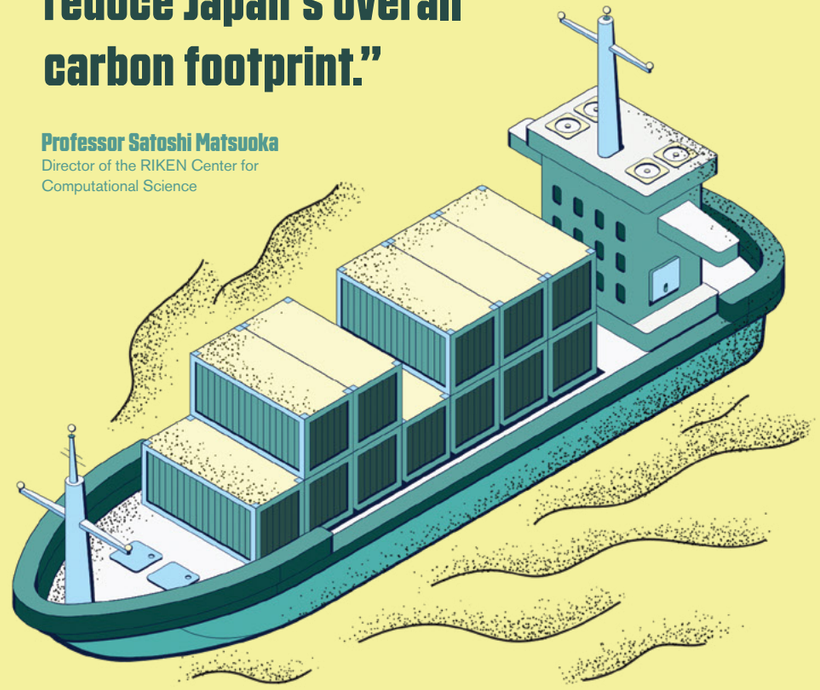

"There are many areas—from power to various industries like food and materials—where Fugaku is contributing," said Matsuoka. "We haven't done a cumulative assessment of just how much potential Fugaku has allowed us to tap into. Maybe that's something we should do, collectively, so we can continue to significantly reduce Japan's overall carbon footprint."

CAPTURING CARBON

After reducing carbon emissions, the other side of the global decarbonization effort is capturing carbon already in the environment, a process called carbon sequestration.

Scientists around the world have begun researching how to trap CO_2 underground, an effort that Fugaku and other supercomputers have assisted in. For this method of carbon sequestration to work, scientists need HPC to simulate millions of scenarios to understand the best way to inject CO_2 into the ground—such as the optimal place and method—and afterwards, to prevent it from escaping.

Another path for carbon sequestration is enhancing the earth's natural systems for balancing CO_2 levels. Researchers at the RIKEN Plant Science Center are looking into using plants both for producing energy and absorbing more CO_2 from the atmosphere using Fugaku.

These efforts are conducted in parallel with research elsewhere in the world. Blue Waters, a petascale supercomputer operated by the University of Illinois, has worked to improve tree mapping efforts for non-forest trees, which function as significant carbon sinks but are less studied than trees found in forests. Meanwhile, the Pawsey Supercomputing Centre in Australia is looking into ways to turn CO_2 itself into fuel.

A SHARED GOAL

For Matsuoka, the race to net zero is not a race among different countries—it's one we are collectively competing in against time. This is why, as part of Fugaku's application-first philosophy, the RIKEN center provides compute cycles to climate and decarbonization projects to be rolled out across Asia.

For instance, one initiative Fugaku will soon be used for is a microalgae-based carbon capture and utilization project for Indonesia under the Japanese government's Science and Technology Research Partnership for Sustainable Development program. Led by Tokyo Institute of Technology's Dr Muhammad Aziz, the project aims to capture CO_2 and use it as a sustainable resource with the use of microalgae. "In addition to its scientific contribution," Aziz told *Supercomputing Asia*, "This study is also expected to help in the reduction of CO_2 and nitrogen oxides from power plants."

Aside from providing access to HPC resources, the RIKEN team also provides training and education to those interested in conducting HPC-enabled research, depending on skill level and project complexity. Matsuoka described this process as one where newbies and seasoned veterans "work their way up the mountain," as Fugaku takes its name from the alternative name for Mount Fuji—Japan's highest peak, but one that Matsuoka says is an easy mountain to climb.

This way, the team behind Fugaku is providing HPC resources for all interested researchers, from beginners to real experts, who want to address our world's biggest crisis. "At the end of the day, it doesn't really matter who solves the problem of carbon," Matsuoka said. "As long as the problem is solved." S

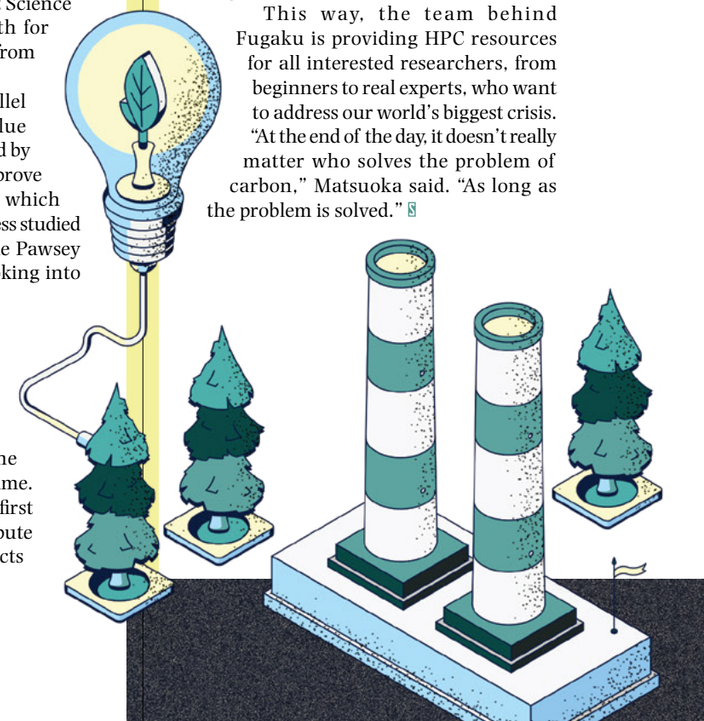

Thoughtful processor design, optimized programming and strong government support are helping to make high performance computing ecosystems more sustainable.

By **Tristan Mañalac**

Illustrations by Shelly Liew / *Supercomputing Asia*

CAPPING CARBON ON ASIA'S SUPERCOMPUTERS

When Pixar and Disney first shared a glimpse of their 2013 animated film *Monsters University*, fans quickly noticed one striking detail: the fur covering Sulley, one of the movie's two protagonists, was extremely realistic.

Millions upon millions of tiny strands of fur swayed convincingly whenever the gentle giant moved his limbs, and even ruffled under a tight shirt, just as you'd expect them to. The animators' secret? A supercomputer that would've ranked among the world's fastest at the time, automatically redrawing every single strand of fur and letting it catch and reflect light with each frame of movement.

Ordinary desktop machines don't have the processing prowess to carry out this type of animation—in fact, even the higher-end versions would have had trouble with it. But animators revealed that Sulley and other monsters in the film, along with every texture, shading and frame, owe their crisp, vivid existence to high-performance computing (HPC).

Able to handle billions of calculations easily, this is the same type of technology that is being leveraged to predict tsunamis, supercharge healthcare innovation and study the origin of super massive black holes. Combining powerful processors, sophisticated software and other cutting-edge computing technologies, HPC employs thousands of computing nodes working simultaneously to complete extremely complex computing tasks much quicker than a regular computer can.

There's just one problem. With great computing power comes great energy liability. Even as HPC systems are helping solve some of the most pressing problems in society in the decade since *Monsters University*, they pose another problem: their massive carbon footprint.

To balance computing power and sustainability, supercomputers across Asia are increasingly being designed with more energy-efficient processors and programming. Meanwhile, governments are waking up to the need for more sustainable energy sources and policies as they shape their growing HPC ecosystems.

HPC'S WORRYING POWER PROBLEM

One of the major drivers of HPC carbon emissions is their steep energy demands. After all, there needs to be a reliable and robust stream of energy to support such intense computing power. The HPC system behind *Monsters University*, for example, comprised 2,000 computers totalling 24,000 cores. Despite this computing power, the movie still took over 100 million CPU hours to fully render. All the while, Pixar's power bill kept racking up.

Moreover, the Frontier system, the world's most powerful supercomputer as of November 2022, needs more than 20 MW of power for its over 8.7 million cores—enough to supply some 52,600 houeholds in Singapore for a month. All told, just powering the world's top 500 supercomputers pumps around two million metric tons of carbon dioxide per year, equivalent to approximately 285,000 households.

The Frontier system, the world's most powerful supercomputer as of November 2022, needs more than 20 MW of power for its over 8.7 million cores—enough to supply some 52,600 households in Singapore for a month.

Plus, any honest accounting of the environmental toll of HPC systems should take stock of the entire ecosystem of technologies that support it. After all, the computing machines themselves form only one, albeit central, part of the equation.

The majority of the energy that flows into supercomputers is dissipated as heat. To manage temperatures and ensure that the machines continue to work properly, computing facilities employ elaborate cooling mechanisms, which themselves often consume a lot of power.

Another peripheral source of carbon emissions in HPC systems is data. The International Energy Agency estimated that in 2021, data centers worldwide used some 220 to 321 TWh of energy—enough to eclipse the consumption of some countries. Given the world's growing reliance on HPC systems, Professor Tan Tin Wee, chief executive of the National Supercomputing Centre (NSCC) Singapore, predicted that as much as 10 percent of the world's energy consumption will come from data center operations in the future. "Energy consumption will be a huge problem," Tan told *Supercomputing Asia*.

THOROUGH THOUGHT AND A POWERFUL HEART

A major solution for the steep energy costs of HPC systems is to maximize computing energy efficiency, explained Professor Satoshi Matsuoka, director of the RIKEN Center for Computational Science, in an interview with *Supercomputing Asia*. The goal, he said, should be to keep power consumption at the lowest possible level while also finding ways to achieve better performance.

RIKEN is home to the Fugaku supercomputer, developed by Japanese company Fujitsu. Since Fugaku debuted in 2020, it has consistently led the TOP500 list of the world's fastest supercomputers. Though it was dethroned by Frontier in June 2022, Fugaku remains a solid contender for the world's most powerful—and energy-efficient—supercomputer, particularly when looking at its real-use conditions.

According to Matsuoka, much of what underpins Fugaku's power is thoughtful, purposeful design. "First, we had to design it efficiently," he said, noting that because they knew that the supercomputer would be used for sustainability research, they specifically built its parts to achieve peak computing performance while doing away with other extraneous functions. "The machine was built with a mindset to save power."

The heart of Fugaku—and largely responsible for its extreme energy efficiency—is the A64FX processor, which was also developed by Fujitsu.

A single A64FX chip contains 48 computing cores divided across four core memory groups (CMG). Each CMG can also contain up to one additional core each, which functions as an assistant. In processor parlance, a core is a small processing unit that can perform computing tasks independently of other cores. The vast majority of computer users will be well-served by machines that have two or four cores. A64FX dials up its performance by having 48.

Each core of the A64FX has a clock speed of 1.8 to 2.2 Ghz, which means that every single core can complete 1.8 billion to 2.2 billion cycles per second. Some simpler computing tasks can be completed within one cycle, while more complex instructions take multiple cycles. Though a bit simplistic, higher clock speeds typically translate to superior computing performance.

Matsuoka noted that aside from Fugaku's processor, the network itself is highly efficient. Where commercial network cards use up 25 to 30 W per node, Fugaku's ethernet over copper networks use 10 to 20 W per node.

> **"The Fugaku Supercomputer was built with a mindset to save power."**
>
> **Professor Satoshi Matsuoka**
> Director of the RIKEN Center for Computational Science

The design for Fugaku also includes precise power control features for users. Whereas most processors operate by having all compute nodes on or off at the same time, Fugaku can be configured to run only the parts relevant to a certain task. "It contributes to significant savings in terms of power usage," said Matsuoka.

These features, along with other engineering innovations, have allowed the Fugaku supercomputer to break performance and power-saving barriers. Compared with the K computer, an earlier Fujitsu supercomputer that was decommissioned in 2019, Matsuoka estimated that Fugaku is about 70 times more powerful in terms of real-use performance. "But power consumption only went up by maybe 20 to 30 percent," he explained. "Thus, compared to its predecessor, the power efficiency of Fugaku is nearly a factor of 50."

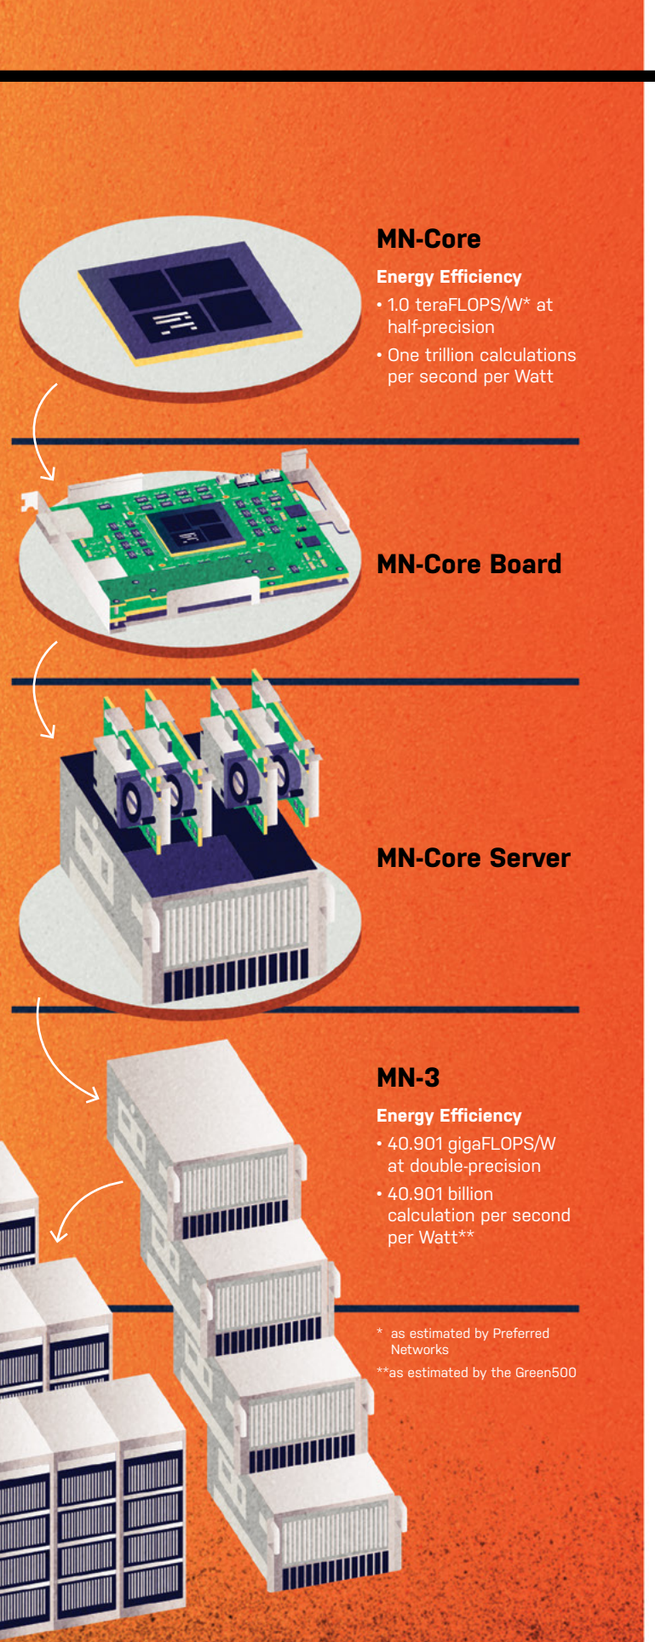

RADICAL SIMPLIFICATION IN SOFTWARE AND CHIP ARCHITECTURE

Peak efficiency is also the objective of MN-3, a supercomputer developed by the Japanese company Preferred Networks, in collaboration with Kobe University.

In fact, despite Fugaku's incredible numbers, MN-3 comfortably eclipses it in terms of energy efficiency. According to Fujitsu's own numbers, for every watt of energy, Fugaku can carry out around 15 billion calculations. With the same amount of energy, MN-3 can perform almost 41 billion—more than double the efficiency.

This impressive statistic has consistently placed the MN-3 among the world's most efficient supercomputers, according to the Green500, a biannual ranking that lists machines in terms of energy efficiency. The MN-3 clinched the top spot in the November 2021, June 2021 and June 2020 lists.

"MN-3 is currently powered by 128 MN-Core processors and 1,536 Intel Xeon CPUs. It consists of 32 nodes with 4 MN-Core processors in each," explained Dr Yusuke Doi, vice president of computing infrastructure at Preferred Networks, in an interview with *Supercomputing Asia*.

However, "the key reason why MN-3 topped the Green500 list three times is precisely that it uses MN-Core, which is specialized for the matrix calculation required for deep learning, instead of GPUs," he added.

MN-Core is an accelerator designed with a hierarchical architecture and comes in a four-die package. Each die has four level-two blocks, which are further divided into eight level-one blocks. In turn, level-one blocks house 16 matrix arithmetic blocks, which themselves contain four processing elements each.

At each level, each block is connected with unique on-chip networks, which can broadcast, aggregate or collect data at every hierarchical level. Different parts of a large dataset can be distributed to different parts of the block, which allows highly efficient processing and computing.

Preferred Networks also employed software optimizations that unlocked the full potential of MN-Core's hardware and helped push MN-3's energy efficiency numbers even higher.

In particular, the company came up with the MN-Core Compiler, a program that translates high-level computer code into another, more machine-friendly language. It was designed with two main goals: to minimize the need for user-side modifications and to maximize MN-Core's features to achieve peak computing performance.

Specifically, the compiler had to figure out the optimal way of mapping out computations to each compute unit in the MN-Core's hierarchical structure. Since the accelerator uses only a single instruction stream, the program also had to ensure a steady flow of data to push performance as close as possible to its theoretical max.

The end result is software that has strong control over hardware and can dictate how calculations will be carried out to achieve maximum efficiency. "In MN-Core, what's conventionally decided and processed within the hardware automatically is exposed to the software side, and the software can manually control details of the computation in the hardware in a 'manual mode' to optimize energy consumption," Doi explained.

"In MN-Core, what's conventionally decided and processed within the hardware automatically is exposed to the software side, and the software can manually control details of the computation in the hardware in a 'manual mode' to optimize energy consumption."

Dr Yusuke Doi
Vice President of Computing Infrastructure at Preferred Networks

This reflects Preferred Networks' core philosophy: realizing hardware's true promise through smart software design. "As long as they are properly controlled by the software, it can unleash silicon's true potential," Doi said.

POWER, POLICY & POLITICAL PRECEDENT

Despite the industry-transforming sustainability efforts of companies like Fujitsu and Preferred Networks, some crucial factors remain beyond the power of private entities.

For instance, in evaluating the carbon emission toll of a supercomputer, looking at how much energy it uses or how efficiently it can carry out calculations isn't enough. It's also important to factor in their country's energy mix. HPC systems in countries powered mostly by renewable energy will be more sustainable than those in territories still reliant on fossil fuels, which is why Matsuoka shared that part of Fugaku's mission is to help Japan develop its offshore wind and solar energy generation. But not every country is able to keep pace.

According to Singapore's Energy Market Authority (EMA), some 95 percent of the country's electricity comes from natural gas. This is the cleanest form of fossil fuel energy—but is nevertheless a carbon-intensive source. EMA estimates that Singapore will continue to rely on natural gas for the foreseeable future, but continues to search for and invest in more sustainable alternatives, like solar energy.

Aside from developing cleaner sources of energy, governments also have the power to shape their countries' HPC ecosystems, coming up with policies that could help them meet consumer and industry demands while also keeping them in line with emission targets.

In Japan, for example, the government has announced substantial subsidies to help data centers make sustainable upgrades to their facilities. The country is also considering concentrating these power-hungry centers in the colder regions of the country, which could help cut back on electricity needs for cooling systems.

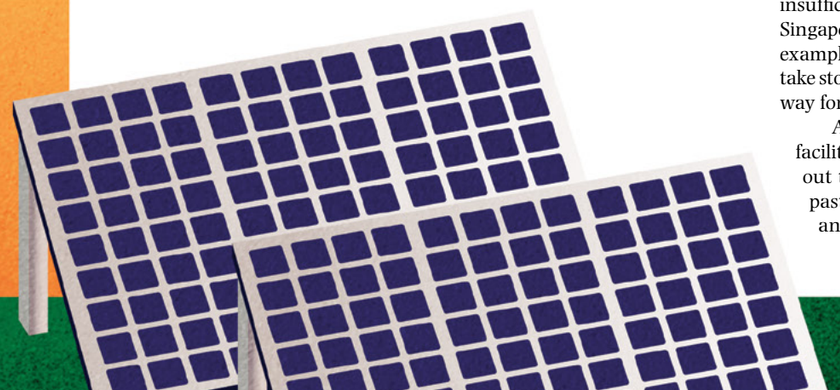

Meanwhile, the Singapore government suspended the approval and construction of new data centers in 2019, pointing to their 350 MW power footprint. The moratorium ended in 2022, and allowed officials to create new guiding principles moving forward.

Under the new rules, only facilities that pass stringent international standards, employ best-in-class energy efficiency technologies and present clear plans to integrate renewables and other innovative energy pathways into operations will be certified. These measures will help Singapore balance the growing need for data centers with the need to respond to the urgent climate crisis.

However, technologies and circumstances are ever-evolving. What may be best-in-class today could be ineffective tomorrow; carbon targets this year could be insufficient the next. In the face of these uncertainties, Singapore has set a good precedent for itself, and a good example for the rest of Asia—push the pause button, take stock of the haves and have-nots and chart a better way forward.

As for NSCC, a government-funded supercomputing facility, chief executive Professor Tan Tin Wee pointed out that their role is to lead by example. Over the past seven years, his team has pioneered cheaper and more efficient cooling techniques that have lowered the energy consumption of their HPC systems—a crucial endeavor for supercomputing in a tropical country. "We can keep trying out new things, which commercial data centers do not have the luxury of doing," Tan explained. "If we can show others that we can do it, then the rest of the community can follow."

"We can keep trying out new things, which commercial data centers do not have the luxury of doing. If we can show others that we can do it, then the rest of the community can follow."

Professor Tan Tin Wee
Chief Executive of the National Supercomputing Centre Singapore

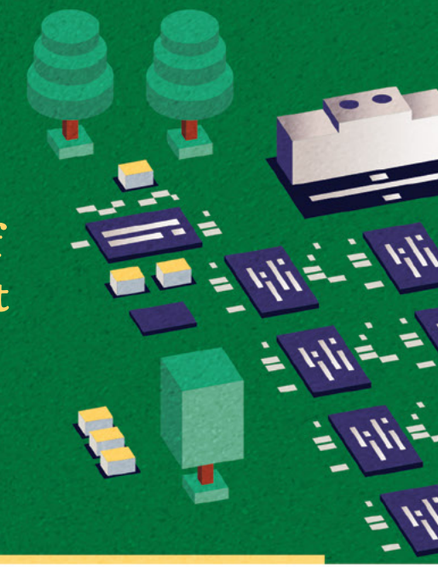

These techniques have been applied to the NSCC's newest supercomputer, the ASPIRE 2A. Designed based on lessons from ASPIRE 1, the ASPIRE 2A has a PUE—or power usage effectiveness, a metric used for measuring a data center's energy efficiency—of close to 1.08. Typical data centers in the region have a PUE of 2. And already, these innovations have been getting some much-deserved recognition. The NUS-NSCC i4.0 Data Centre, which houses the ASPIRE 2A, received the Building & Construction Authority (BCA) Platinum Green Mark Award for Data Centres in 2021 and the W.Media Southeast Asia Cloud & Datacenter (DC) Award for Energy Efficient Innovation in 2022.

To keep improving their systems' energy efficiency, the NSCC also runs simulations of their own supercomputers. In this way, Tan said, "Supercomputers are not just a contributor, but a solution to the problem itself."

Whether HPC systems are used to make the most realistic animated monsters or push the bleeding edge of scientific knowledge, it is important to make sure that their emissions fall in line with the planet's sustainability targets. Asia's innovations in processors, programming and policies have shown that this is possible. §

PULLING THE CURTAINS ON CLIMATE CHANGE

From predicting torrential rains to modeling urban climates, supercomputing provides us the necessary tools to navigate the climate crisis.

By **Mitchell Lim**

Asian countries today are at the mercy of extreme weather events spurred by climate change. Typhoon Hagibis in 2019, for example, is the largest to hit Japan since records began, and monsoon rains submerged thousands of homes in Malaysia at the start of 2022.

From Pakistan to the Philippines, climate change has altered the conditions in which all weather occurs: the oceans and air are warmer, the atmosphere is more vaporous, and sea levels are higher. These result in stronger storms and typhoons, as well as more death and destruction in floods and landslides. A recent study found that climate change made Typhoon Hagibis' extreme rainfall 67 percent more likely, and is responsible for US$4 billion worth of the damage left behind.

As world leaders set goals for mitigating climate change, scientists are helping them prepare for the disasters it will cause. In developing countries, accurate early-warning systems are proving key to improving evacuation systems and limiting casualties. The Global Commission on Adaptation estimates that every dollar spent on such systems could prevent up to twenty dollars of losses annually—underscoring the tangible benefits of coordinated early action.

At RIKEN, Japan's largest research institution, weather and climate scientists are making rain predictions for the Tokyo metropolitan area using Fugaku, previously the world's fastest supercomputer. The key to their technique is big data assimilation: harnessing computational prowess to synchronize data between large-scale computer simulations and observational data.

"For the first time, we have developed a highly accurate and precise model that is capable of predicting weather conditions of up to 30 minutes into the future, updated every 30 seconds, with a resolution of 500 meters," said Dr Takemasa Miyoshi, who leads the Data Assimilation Research Team at the RIKEN Centre for Computational Science, in an interview with *Supercomputing Asia*.

Innovations like these are set to improve the preparedness and response systems for natural disasters. And as climate change continues amidst rapid urbanization and other issues like land sinking, high performance computing (HPC) is empowering researchers in countries like Japan and Singapore to find ways to navigate the climate crisis in the decades to come.

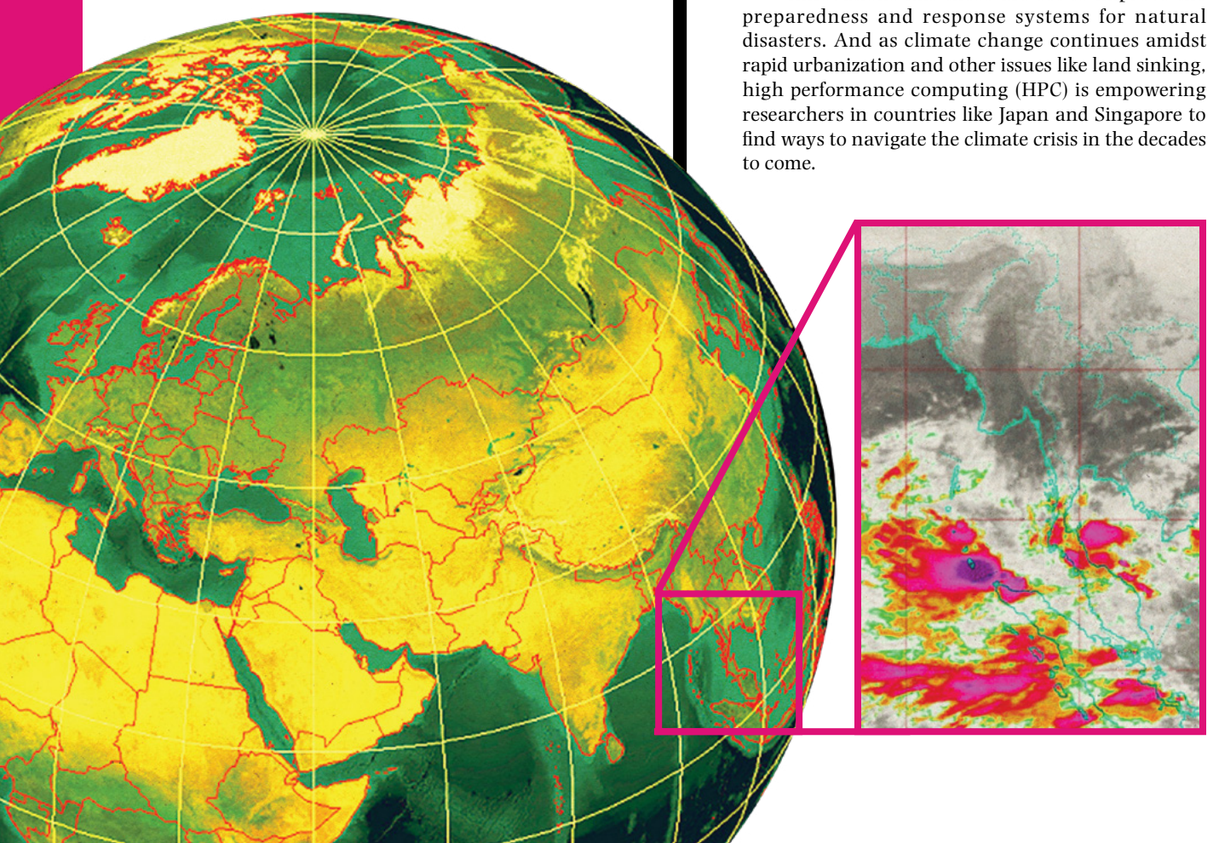

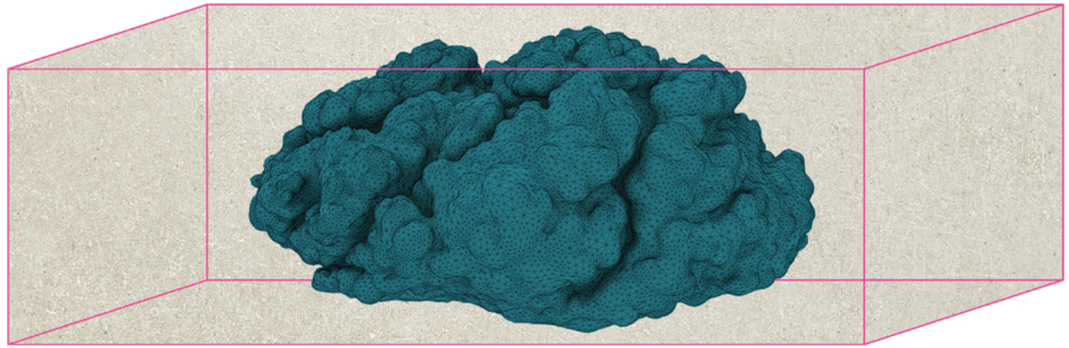

"FOR THE FIRST TIME, WE HAVE DEVELOPED A HIGHLY ACCURATE AND PRECISE MODEL THAT IS CAPABLE OF PREDICTING WEATHER CONDITIONS OF UP TO 30 MINUTES INTO THE FUTURE, UPDATED EVERY 30 SECONDS, WITH A RESOLUTION OF 500 METERS."

DR TAKEMASA MIYOSHI
Team Leader of Data Assimilation Research Team,
RIKEN Center for Computational Science

THE CHALLENGE OF CLOUD SUPERCOMPUTING

"The threat of extreme weather is increasing due to climate change, and it is important to be prepared well," explained Miyoshi. To do this, scientists must develop climate models and weather prediction systems that enhance preparedness with respect to each area's infrastructure, geography and unique weather warning systems.

But this is much easier said than done. It takes a mind-boggling plethora of factors to model the climate, which requires observing the atmosphere, land and oceans with satellites, aircraft and ships.

In fact, modeling clouds is already quite the challenge. At any given moment, more than half of our planet is covered by clouds, which both reflect solar energy to space and absorb heat from the surface of the Earth—influencing weather in different and sometimes contradictory ways.

For instance, surface temperatures can be affected by the height of overhead clouds, how wispy or puffy they are, and how well they reflect sunlight. To study what is going on up above, researchers may even have to probe the deep seas for clues, as ocean microbes have also been found to spit out biological particles that can seed clouds and shape weather patterns across the globe. As ocean temperatures rise and more greenhouse gasses make their way to the atmosphere, these patterns of cloud formation are changing, too.

Clouds play a crucial role in weather and climate, and so there is a lot riding on getting this one set of factors right—an inaccurate computational understanding of clouds leads to an inaccurate climate model and, in turn, inaccurate conclusions and ill-informed policies.

"HPC IS ESSENTIAL TO PERFORM THE TRILLIONS OF CALCULATIONS REQUIRED TO PREDICT THE IMPACT OF CHANGING GREENHOUSE GAS LEVELS ON THE CLIMATE SYSTEM, WHICH INCLUDE SEA LEVEL RISE, RAINFALL AND TEMPERATURE CHANGES."

PROFESSOR DALE BARKER
Director of the Centre for Climate Research Singapore

Clouds must also be considered on top of other factors, like oceans and urban infrastructure. Accounting for the chaos and complexity of all the elements in flux within our planet's systems, today's climate models draw so much brute computing strength that whole supercomputers are designed and built to serve this purpose.

"HPC is essential to perform the trillions of calculations required to predict the impact of changing greenhouse gas levels on the climate system, which include sea level rise, rainfall and temperature changes," said Dale Barker, director of the Centre for Climate Research Singapore (CCRS), in an interview with *Supercomputing Asia*.

TAKING THE LONG VIEW

Apart from monitoring and predicting extreme weather events, supercomputers are also being deployed to address disasters that unravel gradually.

Land sinking, for example, is a particularly pressing problem for Singapore, intensified by the rapid construction of urban infrastructure. According to research led by Nanyang Technological University, Southeast Asia's coastal cities are sinking the fastest, which greatly heightens the risk of coastal flooding when harsh weather strikes. With about a third of its land surface less than five meters above sea level, Singapore is highly vulnerable if no appropriate countermeasures are taken.

Collaborating with the National Supercomputing Center (NSCC) Singapore, scientists at CCRS are tackling this looming hazard by formulating simulations to help understand the complex underlying processes that shape weather and climate in the region. "A better understanding of how these processes will change under future climate scenarios will lead to more reliable climate change projections, which can in turn be used to develop more effective policies for climate change adaptation," Hugh Zhang, deputy director of the Department of Weather Research at CCRS, said in an interview with *Supercomputing Asia*. Zhang also pointed out that because Southeast Asia is considered the "engine room" of the global climate system, this research has the potential to impact policies around the globe.

As part of V3, Singapore's third Climate Change Study, the simulations currently running on NSCC's ASPIRE 2A are expected to culminate in a next-generation climate model tailored for the region by September 2023. Capable of zeroing in on every square of land or sea spanning eight kilometers regionally and two kilometers within Singapore, the model can produce more detailed climate change projections compared to global models, which provide a resolution of about 150 kilometers, for vital information like projected temperatures, rainfall levels and wind speeds.

On the hardware front, the National Environment Agency's Meteorological Service Singapore (MSS) has recently commissioned a new supercomputer named "Utama," built using Hewlett Packard Enterprise's Cray Ex system. Housed in the CCRS building, the Utama has 98 compute nodes and a peak performance of 400 teraFLOPS—computational power that is set to be used to improve weather forecasts and turbocharge climate simulations for Singapore. The new Utama machine will enable upgrades to MSS' 'SINGV' numerical weather prediction system. This includes the demonstration of the new 'uSINGV' capability, designed to capture Singapore's unique urban environment, as well as an upgraded 'cSINGV' configuration that captures the interactions between land, atmosphere and the ocean.

The first two installments of Singapore's Climate Change Study had identified the long-term effects of climate change on the republic, encompassing critical areas such as water resources, biodiversity, public health, infrastructure and food security. With results from the third installment within close sight, supercomputing power—and the research it makes possible—can help shape government policy. It can help identify the appropriate timeframe to implement adaptation measures that could safeguard cities in the region against a warmer climate, rising sea levels and stronger and more frequent storms.

Some measures have already been undertaken as a result of these studies. For instance, the Public Utilities Board, Singapore's national water agency, has rolled out nature-based, engineering solutions to protect the nation's coastlines from rising sea levels, such as putting up mangroves in key areas at risk of high impact waves.

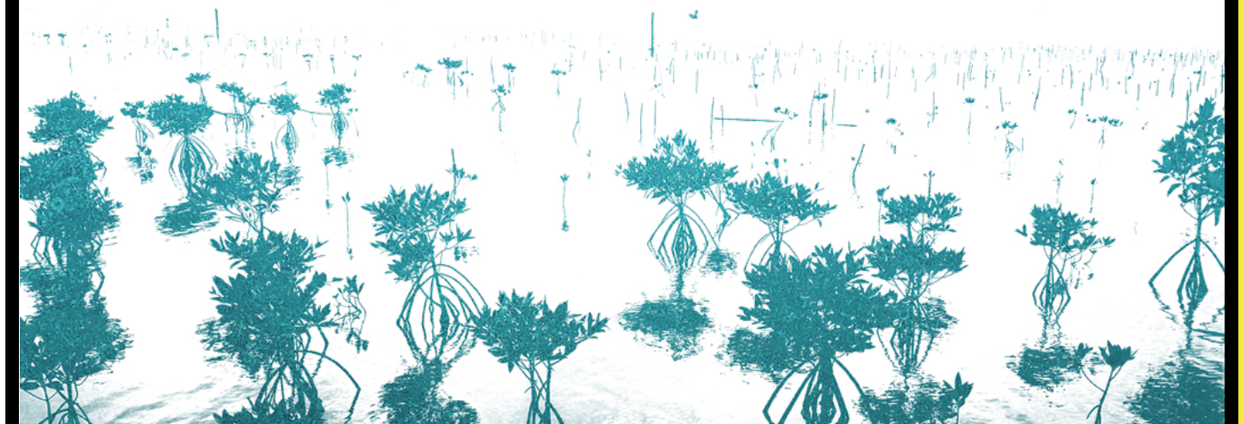

THE ROAD AHEAD

Accurately gauging the risks posed by climate change at a local level is among the key goals of climate modeling. Nevertheless, some uncertainties remain. For instance, researchers are working to zero in on disastrous tipping points, or drastic changes in the climate like the loss of Arctic sea ice or the thawing of carbon-rich permafrost. Once reached, these points can produce climate-destabilizing outcomes that go on to trigger other cascading events, predicted to be irreversible even if global temperatures drop.

Zhang explained that supercomputers will continue to play a key role in advancing weather and climate science. And as capabilities in climate modeling continue to develop, researchers remain hopeful. "Data-assimilation techniques and artificial-intelligence applications could take advantage of advanced hardware, paving the way for future climate research that combines both process and data-driven approaches," said Miyoshi.

"A BETTER UNDERSTANDING OF HOW [WEATHER AND CLIMATE] PROCESSES WILL CHANGE UNDER FUTURE CLIMATE SCENARIOS WILL LEAD TO MORE RELIABLE CLIMATE CHANGE PROJECTIONS, WHICH CAN IN TURN BE USED TO DEVELOP MORE EFFECTIVE POLICIES FOR CLIMATE CHANGE ADAPTATION."

HUGH ZHANG
Deputy Director of the Department of Weather Research at the Centre for Climate Research Singapore

FOSTERING THE NEXT GENERATION OF HPC LEADERS IN ASEAN

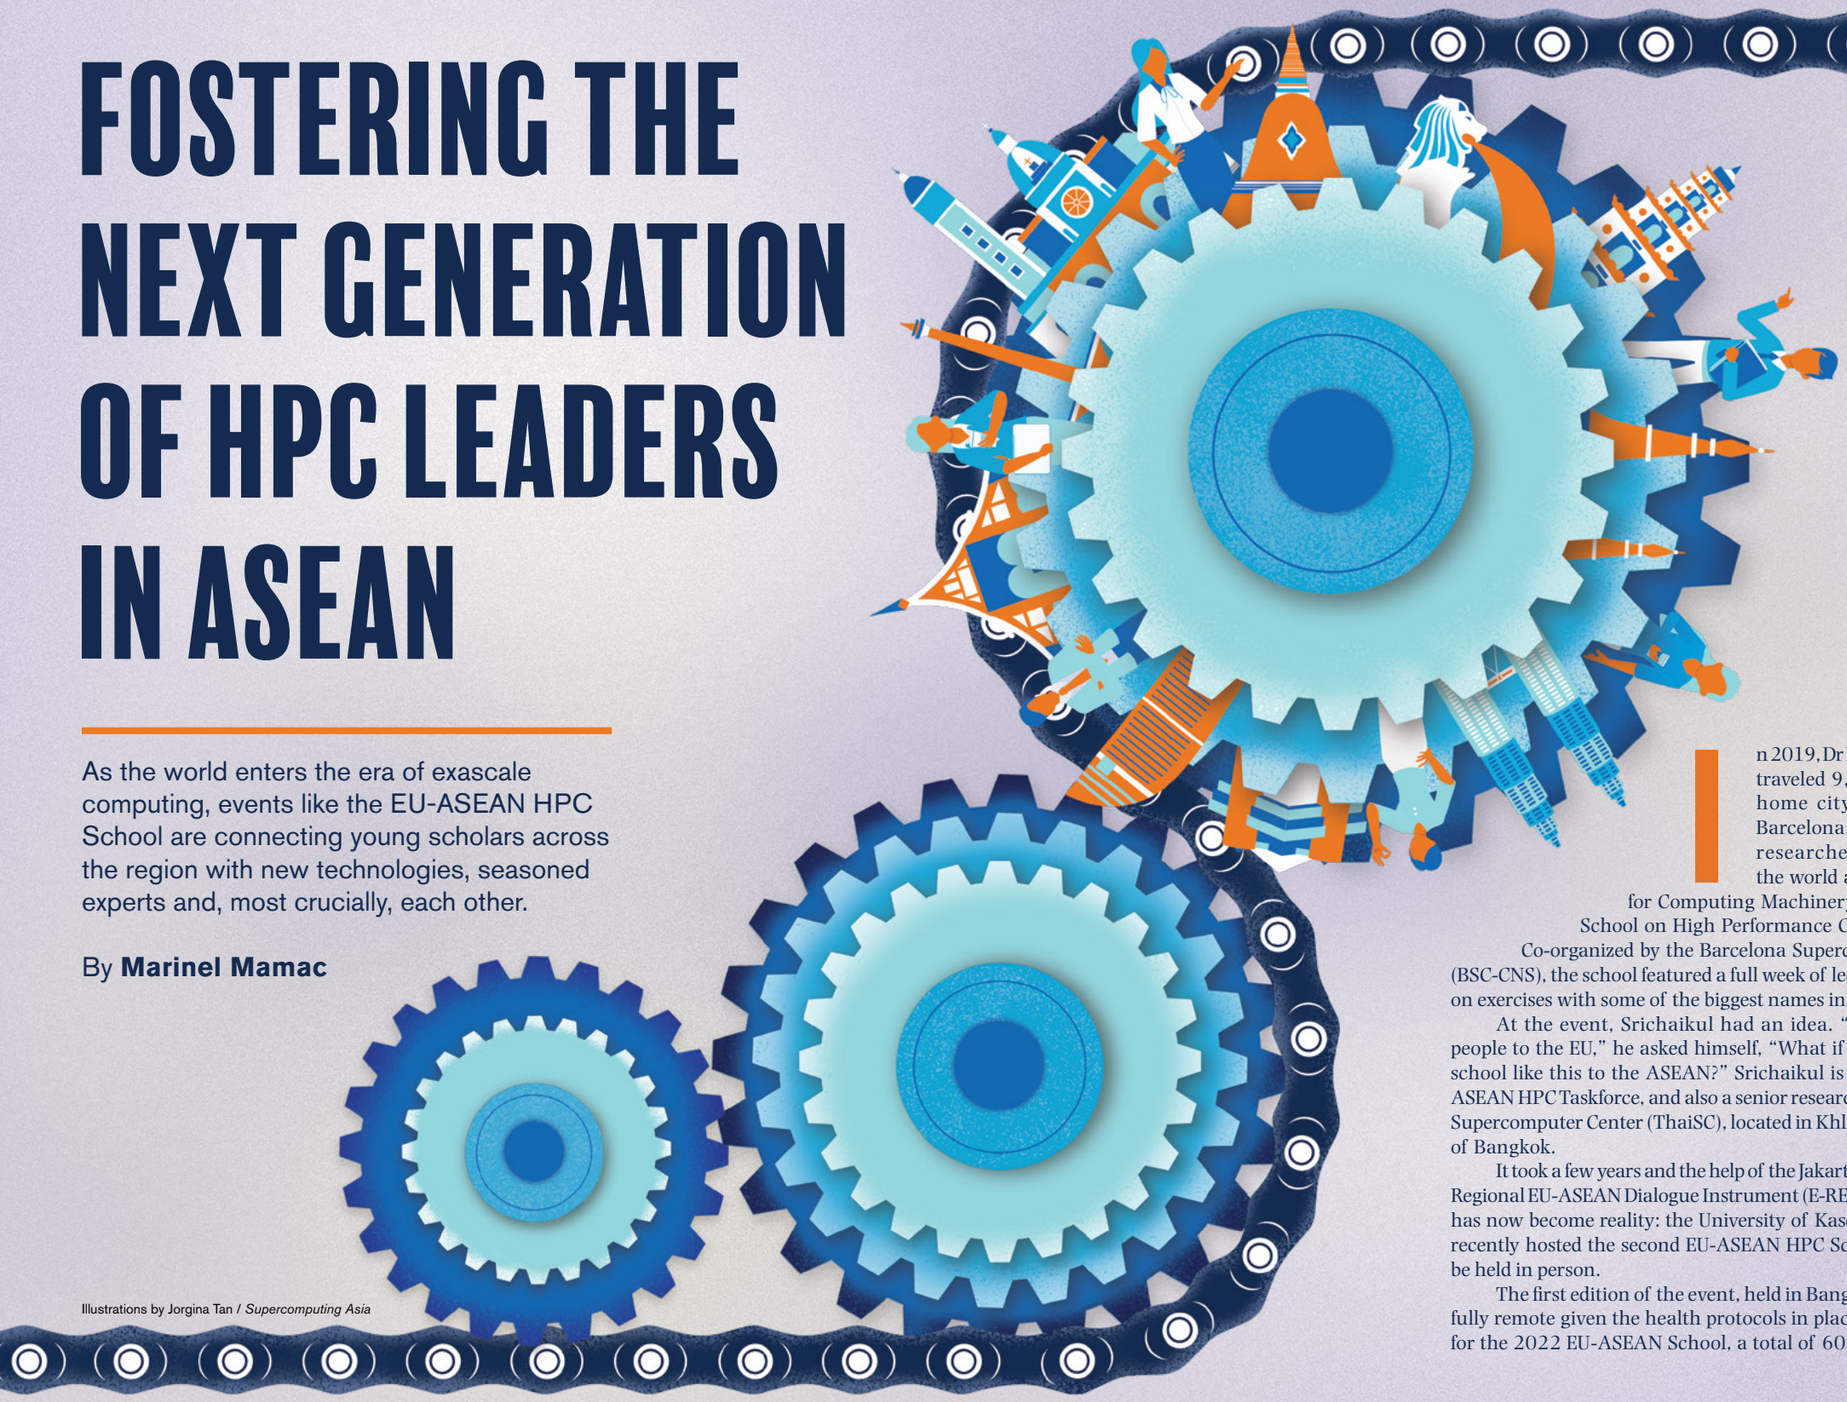

As the world enters the era of exascale computing, events like the EU-ASEAN HPC School are connecting young scholars across the region with new technologies, seasoned experts and, most crucially, each other.

By **Marinel Mamac**

Illustrations by Jorgina Tan / *Supercomputing Asia*

In 2019, Dr Piyawut Srichaikul traveled 9,673 km from his home city of Bangkok to Barcelona, where he joined researchers from around the world at the Association for Computing Machinery (ACM) Summer School on High Performance Computing (HPC).

Co-organized by the Barcelona Supercomputing Center (BSC-CNS), the school featured a full week of lectures and hands-on exercises with some of the biggest names in HPC.

At the event, Srichaikul had an idea. "Instead of flying people to the EU," he asked himself, "What if we could bring a school like this to the ASEAN?" Srichaikul is a co-chair of the ASEAN HPC Taskforce, and also a senior researcher at the NSTDA Supercomputer Center (ThaiSC), located in Khlong Nueng, north of Bangkok.

It took a few years and the help of the Jakarta-based Enhanced Regional EU-ASEAN Dialogue Instrument (E-READI), but this idea has now become reality: the University of Kasetsart in Bangkok recently hosted the second EU-ASEAN HPC School—the first to be held in person.

The first edition of the event, held in Bangkok in 2021, was fully remote given the health protocols in place at the time. But for the 2022 EU-ASEAN School, a total of 60 students selected from a pool of 300 applications from all 10 ASEAN member states were able to attend the event in person. The School ran from December 5 to 10, 2022, and is one of a growing number of programs and activities in the region encouraging young scholars to learn more about HPC-enabled research.

CREATING VALUABLE CONNECTIONS

Srichaikul explained that the EU-ASEAN HPC School has two goals. The first is to create a talent pipeline for the HPC industry. "We want to grow local talent and promote capacity-building among ASEAN member states," he told *Supercomputing Asia* in an interview. "We want students to learn and translate that learning to real benefits."

The second, more long-term goal is promoting regional collaboration. "You cannot expect every nation to have the same technology," Srichaikul said, explaining that while countries like Singapore and Thailand have national-scale HPCs in place, others are still laying out roadmaps to do the same. "ASEAN countries have economic gaps, and so we have to ask: How can we make HPCs into a regional resource so that our scientists can tap into this computing power?"

Dr Fabrizio Gagliardi, director of the EU-ASEAN HPC School, shares this view. "ASEAN is a little bit like Europe," Gagliardi pointed out in an interview with *Supercomputing Asia*. "There are big differences, say, between countries like Germany and Romania. Their economies are completely different, the way it is with Singapore and Vietnam. And so, HPC schools are meant to be helping those who may be running a little behind."

These two goals are why Gagliardi and his team designed the school to not only have sessions on the most powerful HPC systems in the region—such as Singapore's ASPIRE 2A and Thailand's LANTA, aside from Japan's Fugaku and Finland's LUMI—the program also included plenty of opportunities for participants to network with each other and with the school's esteemed guests.

For the 2022 EU-ASEAN HPC School, the guests of honor included Satoshi Matsuoka, director of the RIKEN Center for Computational Science; Anders Jensen, executive director of EuroHPC; Jack Dongarra, Turing Laureate 2021 and author of the The LINPACK Benchmark that serves as the basis for the TOP500 list; and Supa Hannongbua, president of the Chemical Society of Thailand.

For his part, Gagliardi is no stranger to fostering younger researchers and networks. As a senior strategy advisor at the BSC-CNS, he has overseen several HPC summer schools in the EU and was part of the team that brought Srichaikul to Barcelona back in 2019.

Now Srichaikul's colleague in organizing the EU-ASEAN HPC School, Gagliardi shared that it all boils down to developing the next generation of HPC scientists. "I think it's part of any scientist's mission. We work hard all our lives—which are relatively short—and we reach a level of knowledge we don't want to be wasted, and so we want to train the younger generation," he said. "We want to ensure that there are people behind us to continue the work." The same can be said for the guest speakers and lecturers, who came all the way to Bangkok in 2022 as volunteers.

By connecting students to state-of-the-art HPC technology, seasoned experts and each other, the EU-ASEAN HPC School hopes to develop participants' potential not just for their individual research projects, but for the region in general.

Dr Marieanne Leung, an atmospheric scientist who received the Best Student Award in the 2021 EU-ASEAN HPC School, agrees. "My hope is for an open-source or shared HPC facility to be available in the ASEAN region," she shared in an interview with *Supercomputing Asia*. "These facilities are expensive, and not all countries or organizations can afford it." With this shared HPC environment, Leung hopes to pool regional expertise to tackle big issues like climate change.

"We have good people," Srichaikul confidently said. "They know what they want to do. It's just a matter of giving them the opportunity [to achieve it]."

"THE TIME IS NOW, AND THE OPPORTUNITY IS HERE."

For Gagliardi, there is no better time to grow HPC talent in the ASEAN.

"It has always been very exciting," said Gagliardi. "But now, there are so many transformations at the level of electronic components and the ways you put them together to interconnect memory hierarchies. And then on the software side, there are new ways to program all these large machines."

With that in mind, he encourages young people to seek out different HPC programs. "The time is now, and the opportunity is here, regardless of whether you get into the EU-ASEAN HPC School or not."

And he is right: aside from the school, there are webinars, scholarships and competitions for young scientists interested in learning more about HPC.

In 2022, for example, the National Supercomputing Centre (NSCC) Singapore organized two competitions. The first was the 2022 APAC HPC-AI Competition, for which the NSCC teamed up with the HPC-AI Advisory Council and the National Computational Infrastructure (NCI) Australia. Open to teams from across Asia and the Pacific, the competition provided training in artificial intelligence, HPC foundations, UCX programming and Quantum Espresso, among others. Winners were announced in Dallas, Texas, in November, and will receive their awards at the Supercomputing Asia 2023 conference in Singapore this March.

The second landmark competition was the Inaugural HPC Innovation Challenge for the Environment, which was supported by GeoWorks, SGTech and SGInnovate. Open to local enterprises and college students, the program challenged its participants to think of HPC-enabled solutions that can support data-centric approaches to manage the environment, reduce one's carbon footprint, create better urban environments and build climate resilience. A total of 10 teams across the student and open categories were shortlisted for the solution development phase, where they were given access to the ASPIRE 2A as well as mentorship from subject matter experts.

"The environmental solutions developed by the winning teams in the HPC Innovation Challenge are just the beginning," said Professor Tan Tin Wee, chief executive of NSCC. "Singapore is brimming with interest and talent for HPC and interest in the rest of the region is rising rapidly. Through events like this, we are very happy to empower young people in diving into this new era of HPC research."

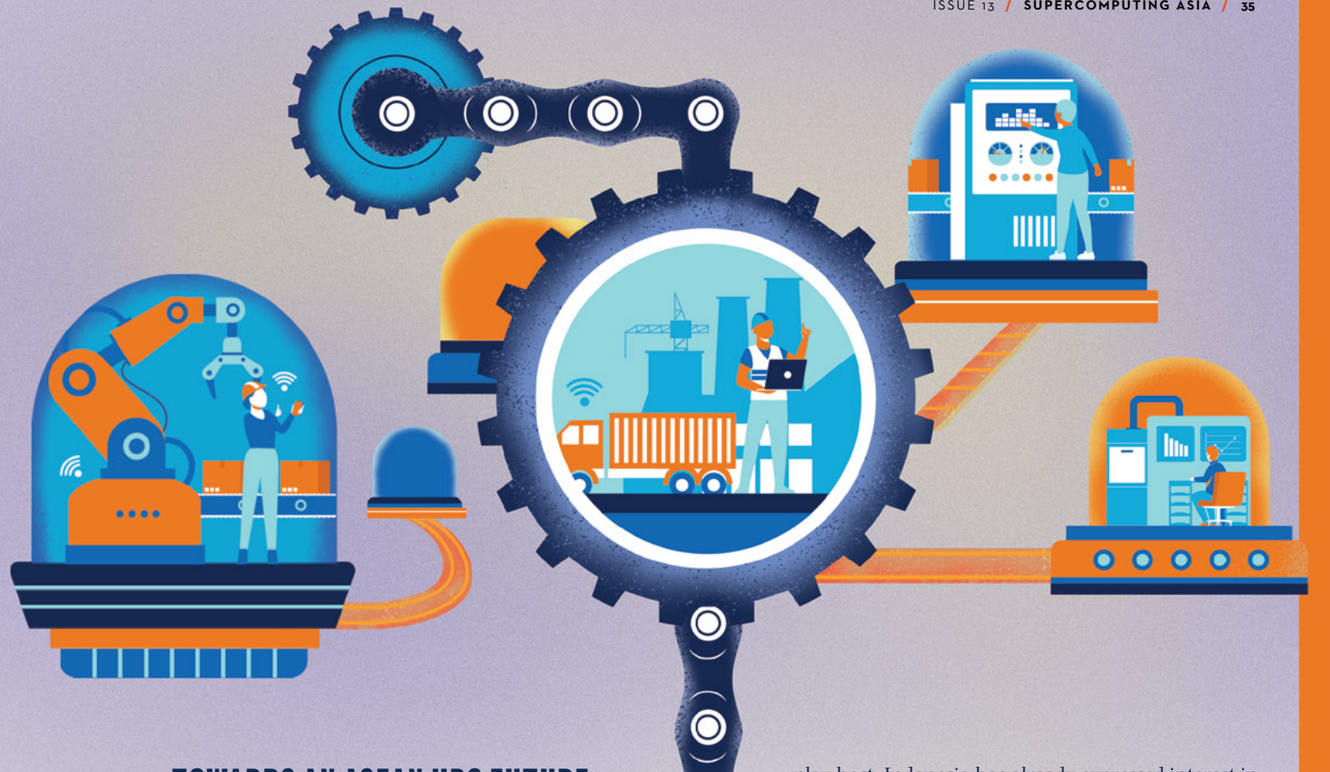

TOWARDS AN ASEAN HPC FUTURE

Taken together, these activities are providing a glimpse of a thriving ASEAN HPC environment—a future that isn't too far, and one that promises exciting findings in a wide range of research areas.

Dr Nikman Adli bin Nor Hashim, a Malaysian genomic scientist who participated in the EU-ASEAN HPC School last December, is looking to build on his new experiences in HPC technology for omics research. "It is undeniable that bioinformatics and computational tools are crucial in today's research, since we are dealing with heavy computational analyses," he told *Supercomputing Asia* in an interview. He hopes to share his learnings with his colleagues to further improve research on Southeast Asian populations.

Leung, for her part, is currently using HPC to study atmospheric processes in Peninsular Malaysia. She was invited again to participate in the second EU-ASEAN HPC School, this time in person. Leung described the experience as "beyond awe-inspiring," and hopes to use her research to contribute to climate-resilient development.

Gagliardi plans to make the EU-ASEAN HPC School a permanent annual event, with countries taking turns to play host. Indonesia has already expressed interest in hosting next year's School.

At the end of the day, Srichaikul says that efforts like the EU-ASEAN HPC School are less about expensive pieces of hardware and infrastructure. "It's really about who is using it and for what benefit," he said. "These research applications are what bring value to ASEAN communities, today and in the future." S

"WE HAVE GOOD PEOPLE. THEY KNOW WHAT THEY WANT TO DO. IT'S JUST A MATTER OF GIVING THEM THE OPPORTUNITY [TO ACHIEVE IT]."

Dr Piyawut Srichaikul
Co-chair of the ASEAN HPC Taskforce and Director of the Thai Supercomputer Center

Business Bytes

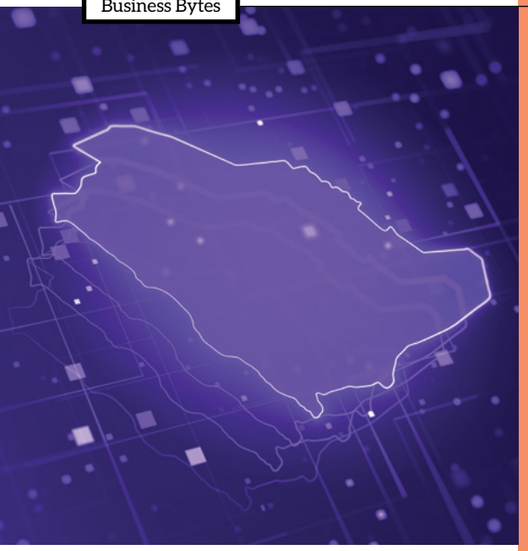

INTRODUCING SHAHEEN III, SAUDI ARABIA'S FASTEST HPC YET

After launching Shaheen I in 2009 and Shaheen II in 2015, Saudi Arabia's King Abdullah University of Science and Technology (KAUST) will debut Shaheen III, or what will be the Middle East's most powerful supercomputer, in 2023.

Set to be 20 times faster than its predecessor, the Shaheen III is expected to deliver 100 petaFLOPS of performance with its 25-cabinet HPE Cray EX4000 system and over 2,800 NVIDIA Grace Hopper superchips. The system is built by Hewlett Packard Enterprise (HPE) and includes AMD's Genoa CPUs, HPE Slingshot networking cables and a Cray ClusterStor E1000 storage system. Its 100-petaFLOPS performance would put the Shaheen III behind the United States Department of Energy's Summit and ahead of its Sierra Supercomputer—currently ranked fifth and sixth, respectively, on the TOP500 list.

Tony F. Chan, president of KAUST, has said that the new supercomputer will help university scientists and partner institutions conduct research on a much larger scale to address some of the region's most significant issues, such as clean combustion, Red Sea ecosystems, the Arabian tectonic plate and climate change.

FUGAKU'S NEXT MOVE

Fugaku, Japan's top supercomputer, debuted in 2020 after six years of development. In the years since, Fugaku has been instrumental in solving research problems related to COVID-19, climate change and other crucial issues of our time, while setting world records in terms of speed and efficiency.

This fiscal year, Japan's Ministry of Education, Culture, Sports, Science and Technology (MEXT) is setting its sights forward with Fugaku-Next. With an overall budget of JPY 450 million (US$3.07 million), MEXT has begun considering the country's next supercomputer with feasibility studies that are set to run until March 2024.

The studies will investigate a new supercomputer's feasibility and profitability, among others, and will be undertaken by four research teams selected from seven applicants: Riken Center for Computational Science, Kobe University, Keio University and Tokyo University. Among the teams in these four institutions are some of Japan's leading scientists in the field, such as Riken's Masaaki Kondo, Kobe University's Junichiro Makino, Keio University's Hideharu Amano, and Tokyo University's Toshihiro Hanawa.

The results of the feasibility study will determine if development will begin in 2024.

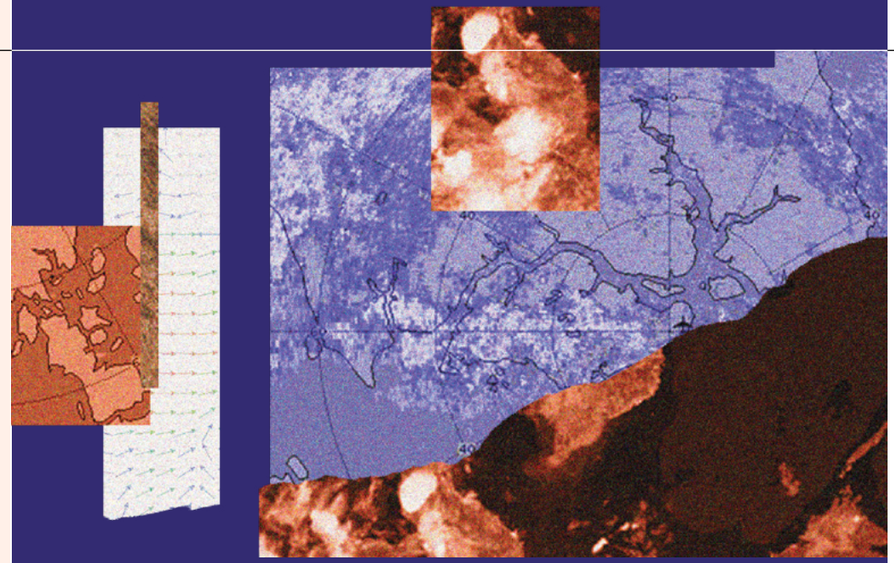

UTAMA, SINGAPORE'S SUPERCOMPUTER FOR CLIMATE RESEARCH, AWAKENS

The National Environment Agency's Centre for Climate Research Singapore (CCRS) has commissioned a new supercomputer. Its name, "Utama," is a nod to the country's heritage: it means "most important" in Malay, and it also references Sang Nila Utama, the Srivijayan prince who founded the kingdom of Singapura in 1299.

With its ability to perform 400 trillion computations per second, Utama is set to boost the agency's weather and climate research—helping Singapore become safer and more climate-resilient.

One area that Utama will aid with is AI-powered weather forecasting. Because of Singapore's tiny size, existing global climate models struggle to make accurate predictions for the development of small and transient weather systems in the area. With Utama, the CCRS aims to solve this issue, using its computing power to help the country better prepare for thunderstorms, monsoons and squalls.

On the other hand, the new supercomputer is also set to help CCRS better understand weather processes and climate change in the region. The machine will be used to develop an ultra-high resolution urban modeling system to accurately capture the urban environment of Singapore, and predict how weather and climate patterns might affect it.

A NEW NATIONAL SUPERCOMPUTER FOR KOREA

The Korea Institute of Science and Technology Information (KISTI) is set to unveil South Korea's most powerful supercomputing system yet in 2023. It will be the sixth of its kind since the country first embraced the technology in 1988.

The new supercomputer is expected to be 32 times faster than its predecessor Nurion, which was launched in 2018 and currently has an average usage rate of around 75 percent. Though there is no information yet on the chosen supplier or the parts that would make up the new supercomputer, it is set to have a peak speed of 600 petaFLOPS, which Lee Sik, director general of KISTI's national supercomputing division, predicts will place the system between 8th and 11th in the world. It reportedly has a budget of KRW 292.95 billion (US$204.5 million).

The move is part of South Korea's efforts to promote its science and technology sector. It will be used in the development of batteries and chips, the simulation of next-generation nuclear reactors, the advancement of space technologies, and the development of new drugs.

Super Snapshot

GOING FARTHER, TOGETHER

Quantum computing is a team effort, and these local and international collaborations are helping to lead Singapore into the quantum age.

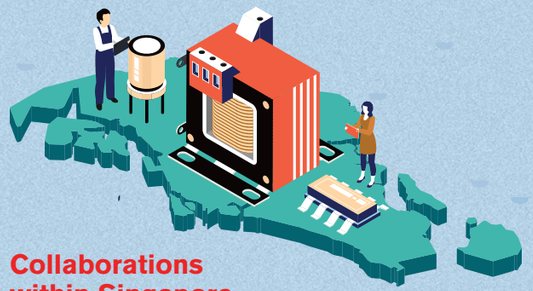

Collaborations within Singapore

To support Singapore's development in quantum computing, quantum-safe communication and quantum device manufacturing, the country's Quantum Engineering Programme has rolled out three national platforms, namely:

- **National Quantum Computing Hub (NQCH)** is developing quantum computing capabilities, pooling expertise from the Centre for Quantum Technologies' (CQT) teams at the National University of Singapore (NUS) and Nanyang Technological University (NTU), as well as the Institute of High Performance Computing (IHPC) at the Agency for Science, Technology and Research (A*STAR). The National Supercomputing Centre (NSCC) Singapore is set to host a quantum computing facility and provide the supercomputing resources required to develop algorithms for quantum computing.
- **National Quantum Fabless Foundry (NQFF)** is hosted at A*STAR's Institute of Materials Research and Engineering (IMRE), and will be supporting micro and nanofabrication techniques for quantum devices and enabling technologies.
- **National Quantum-Safe Network (NQSN)**, led by the CQT and supported by NUS and NTU, is conducting trials of quantum-safe communication technologies nationwide, with the goal of enhancing network security for critical infrastructure.

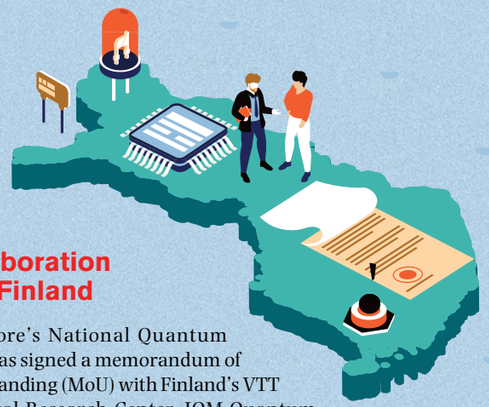

Collaboration with Finland

Singapore's National Quantum Office has signed a memorandum of understanding (MoU) with Finland's VTT Technical Research Center, IQM Quantum Computers and CSC - IT Center for Science to strengthen research cooperation in the area of quantum technology. More specifically, the MoU covers the development of quantum algorithms, applications and hardware components, as well as quantum-accelerated high-performance computing.

Collaboration with France

In the realm of quantum communications, aerospace firm SpeQtral is set to launch SpeQtral-1, its first quantum key distribution (QKD) satellite in 2024. The Singapore-based company is working with the country's Office for Space Technology and Industry as well as Thales, a French aerospace company. SpeQtral-1 is set to be in orbit in 2024.

SpeQtral-1's pilot phase will allow companies to secure their data communications with quantum keys—a crucial step to establishing global quantum-secure networks.These networks are expected to be computationally uncrackable.

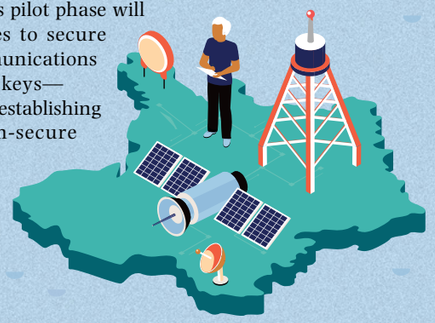

The three national platforms aim to coordinate research organizations and foster public-private collaborations.

Smarter fights global warming with water-cooled servers.

www.lenovo.com/smarter

Smarter technology for all

Lenovo