ASIAN SCIENTIST

Issue 10
July 2021

SUPERCOMPUTING ASIA

SOUTHEAST ASIA'S SUPERCOMPUTING SIGHTS

SETTING A COURSE FOR SUPERCOMPUTING

ACCELERATING HPC ADOPTION ACROSS ASEAN

THE STATE OF SUPERCOMPUTING IN SOUTHEAST ASIA

CONTENTS

Issue 10
July 2021

EDITOR'S NOTE

When it comes to raw supercomputing power in Asia, China and Japan stand out for their impressive rise and deep expertise respectively. But there is more to Asia than just those two countries. Although it may be starting from a lower base, Southeast Asia—with its vibrant young population and dynamic economy—looks set to experience rapid growth.

The region also faces many problems that supercomputers are poised to address, from climate change to emerging infectious diseases. Yet for all its potential and a clear need for high performance computing, efforts to develop Southeast Asia's capabilities have remained relatively disparate—until now. In our cover story (*Accelerating HPC Adoption Across ASEAN,* p. 16), we trace a major collaborative effort designed to improve regional HPC access: the Association of Southeast Asian Nations (ASEAN) HPC shared facility.

While a regional shared facility does not yet exist, individual countries in Southeast Asia already have national and local facilities of their own. Head on over to p. 10 to check out *Southeast Asia's Supercomputing Sights,* a virtual tour of some of the most interesting deployments in the region.

Last but not least, keep your eyes peeled for an exciting development in number formats coming out of Singapore: posit computing (*Posits: Coming Soon to Hardware Near You,* p. 32). Promising to give greater accuracy with fewer bits, posits are gaining momentum and making their way into hardware. You heard it here first!

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

EDITOR-IN-CHIEF
Dr. Rebecca Tan

STAFF WRITERS
Kamila Navarro
Jill Arul
Erinne Ong
Andre Encarnacion
Sheryl Lee

CONTRIBUTORS
Eugene Low
Tim Hornyak

ART DIRECTOR
Lam Oi Keat

SENIOR DESIGNER
Shelly Liew

DESIGNERS
Ajun Chuah
Alexandra Valino

SALES & MARKETING
Samantha Yeap
Jillene Kee

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

THE MUTATIONS BEHIND THE SPREAD OF SARS-COV-2

Aided by the Frontera supercomputer at the Texas Advanced Computing Center, scientists have pinpointed the key mutations that make SARS-CoV-2 variants more contagious. "The UK, South Africa and Brazil variants are more contagious and escape immunity easier than the original virus," said Dr. Victor Padilla-Sanchez, a research scientist at the Catholic University of America. "We need to understand why they are more infectious and, in many cases, more deadly."

Since the viral spike glycoprotein is key for binding to human cells, Padilla-Sanchez computationally modeled the protein's structure to reveal how changes can facilitate infection. For this memory-demanding task, he turned to Frontera's 39 petaFLOPS of computing power and 16 large memory nodes.

In the fast-spreading 'Delta' B.1.1.7 variant, the N501Y mutation altered the spike protein's binding domain so it can latch onto human cells more efficiently. Meanwhile, the 'Beta' 501.V2 variant owes its high transmissibility to the K417N mutation, which strengthens the attractive forces between the spike protein and the human ACE2 receptor. This variant also carries the E484K mutation, aptly called an escape mutant for preventing antibodies from acting on the virus and ultimately letting it slip past the body's immune defenses.

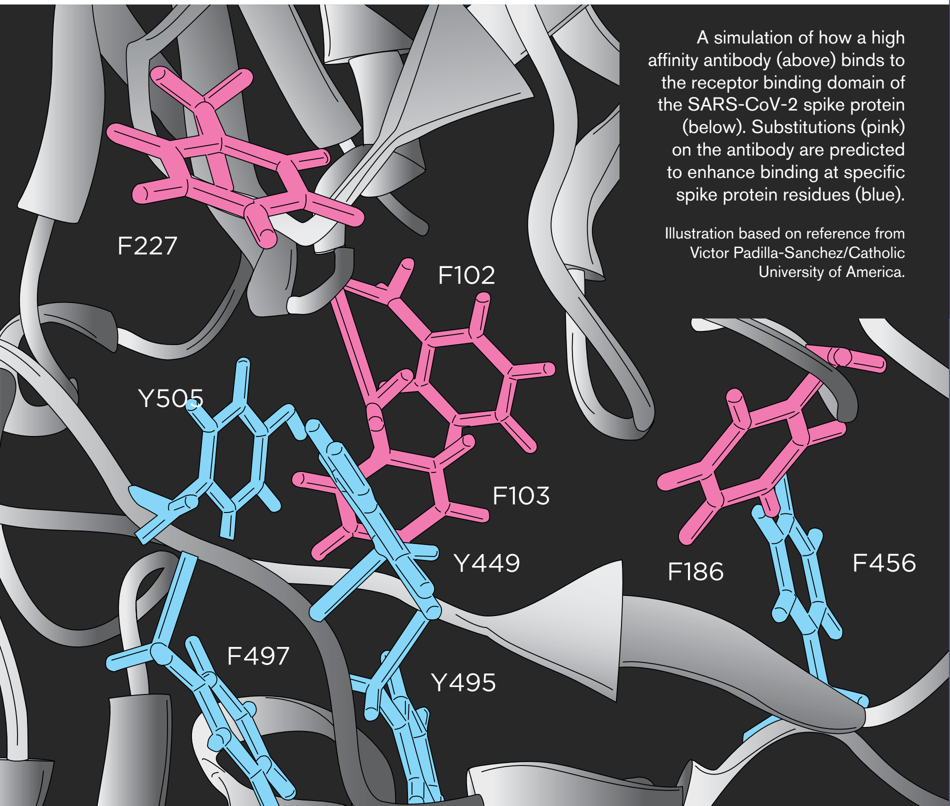

A simulation of how a high affinity antibody (above) binds to the receptor binding domain of the SARS-CoV-2 spike protein (below). Substitutions (pink) on the antibody are predicted to enhance binding at specific spike protein residues (blue).

Illustration based on reference from Victor Padilla-Sanchez/Catholic University of America.

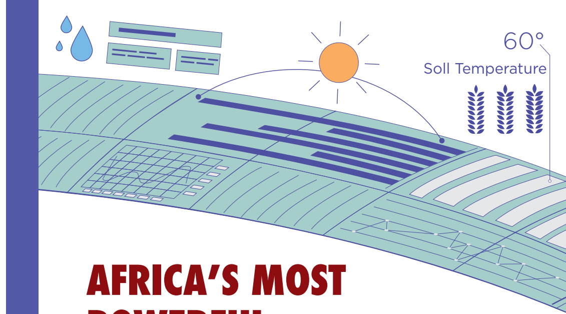

AFRICA'S MOST POWERFUL SUPERCOMPUTER AWAKENS

With a computing capacity of 3.15 petaFLOPS, Toubkal claims the crown as the most powerful supercomputer on the African continent, overtaking the 1.03 petaFLOPS Lengau supercomputer housed at South Africa's Centre for High Performance Computing. Toubkal also stands as the 98th most powerful supercomputer worldwide, marking the African continent's return to the TOP500 since falling off the list in June 2019.

Deployed by the Mohammed VI Polytechnic University (UM6P) in Morocco, Toubkal was manufactured by American multinational corporation Dell EMC and developed in conjunction with the UK's University of Cambridge. The 1,300-node system boasts a main memory capacity of 244 terabytes, delivering excellent performance for all sorts of computing tasks.

These features are already being put to good use at UM6P's new African Supercomputing Center (ASCC), which was inaugurated alongside Toubkal's deployment at the campus. Ongoing workloads include analyzing satellite data for better agricultural management and integrating meteorological data in renewable energy plans. With the ASCC advancing Africa's supercomputing research, researchers and entrepreneurs across the continent can look forward to having their computing needs answered as they take full advantage of the site's resources.

INDIA'S SUPERCOMPUTING STRATEGY GAINS GROUND

By the end of the National Supercomputing Mission's second phase in September 2021, India's supercomputing capacity is set to reach 16 petaFLOPS. The country's goal to become a global supercomputing leader is set to pick up speed in Phase 3, with the ambition of further raising capacity to a whopping 45 petaFLOPS.

By plugging into India's national supercomputing infrastructure, 75 institutions will gain ready access to high performance computing (HPC) facilities for their research and innovation efforts. One major milestone in this nationwide endeavor is the PARAM Siddhi-AI supercomputer, which was inaugurated in November 2020 at the Centre for Development of Advanced Computing. Ranked 63rd fastest globally, PARAM Siddhi-AI is no slouch—combining HPC and artificial intelligence (AI) technologies for a theoretical peak power of 5.26 petaFLOPS.

Meanwhile, a more powerful 20 petaFLOPS system is being readied for deployment as a national facility, complemented by intensive training efforts to develop a ready pool of HPC experts. With supercomputers already pushing limits around the world, India is poised to be transformed by high performance computing, in areas ranging from data-driven healthcare services to cost-efficient sustainability solutions.

PREDICTING TSUNAMIS IN REAL TIME

Ten years after the fateful 2011 Tohoku earthquake and tsunami that claimed nearly 20,000 lives, Japanese researchers have trained a highly accurate AI model to predict tsunami flooding in near real-time for disaster mitigation.

In this joint project, the International Research Institute of Disaster Science at Tohoku University, the Earthquake Research Institute at the University of Tokyo, and Fujitsu Laboratories ran simulations of 20,000 possible tsunami scenarios on Fugaku, the world's fastest supercomputer.

Thanks to Fugaku's unparalleled computing performance, the datasets not only were rendered in high resolution, but were also directly fed into the AI model's deep learning algorithm. The resulting model is able to estimate detailed flooding conditions based on waveform data collected offshore, buying valuable time for people on the ground to react.

Although the model was trained on Fugaku, it can be run in mere seconds on ordinary computers—a rare feat for a supercomputing project—making it highly practical and easy to implement across devices. Moving forward, the researchers will add training data covering a larger land area, gearing the system towards applications in data-driven disaster management.

WEATHER REPORT: FEWER, BUT FIERCER, TROPICAL CYCLONES

Worsening global warming will intensify tropical cyclones making landfall, according to a study published in *Science Advances* by researchers from the IBS Center for Climate Physics (ICCP) at Pusan National University in South Korea.

The fluctuating atmospheric and oceanic conditions of a brewing storm make it difficult to model in extensive detail. This challenge made previous generations of climate models prone to errors, largely due to confusing feedback from ocean temperature data. With its ultrahigh resolution, ICCP's Earth System Model accurately captured atmospheric and oceanic interactions, leading to more realistic simulations of global climate patterns.

To attain this unprecedented level of precision, their model relied on 1.43-petaFLOPS Aleph, ICCP's in-house supercomputer, to sift through global warming data and tropical cyclone characteristics. After 13 months of running simulations, the researchers found that strong tropical cyclones will become less frequent yet carry much more destructive potential, increasing the risk for extreme coastal flooding.

"By representing coastal processes more accurately than ever before in a global model, we now have a much higher confidence in these robust model projections in particular for landfalling tropical cyclones," said ICCP director and study co-author Professor Axel Timmermann.

Digital Dispatch

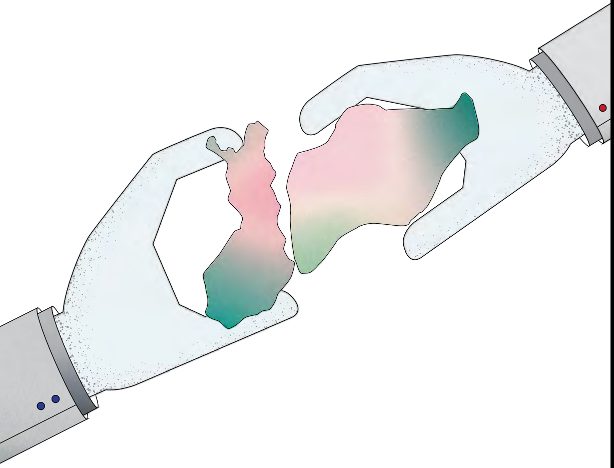

HAND IN HAND FOR GLOBE-SPANNING HPC

Imagine a high-speed, high-bandwidth fiber optic link between two nations halfway across the world from one another. Supercomputing centers in Singapore and Finland may soon turn this vision into reality, as part of a new memorandum of understanding (MOU) for pursuing world-leading HPC projects.

Announced during the Supercomputing Asia 2021 virtual conference, this cross-border collaboration is set to involve Singapore's National Supercomputing Centre (NSCC), Singapore Advanced Research and Education Network, National University of Singapore's quantum engineering program, and Finland's CSC-IT Center for Science.

Alongside a fiber optic network, the four teams are jointly developing quantum-based solutions for more secure data transmission over long distances. Sustainability is also on the agenda for the NSCC and CSC. Given the high power consumption of data centers, they plan to integrate renewable energy sources in designing energy-efficient data centers following models built by Finnish researchers.

"The MOU is a reflection of the resolve by the HPC community to continue cooperation as a tool in fighting global challenges like the current pandemic and in better preparing us to face the challenges to come," said NSCC chief executive Associate Professor Tan Tin Wee.

WHAT'S UP!

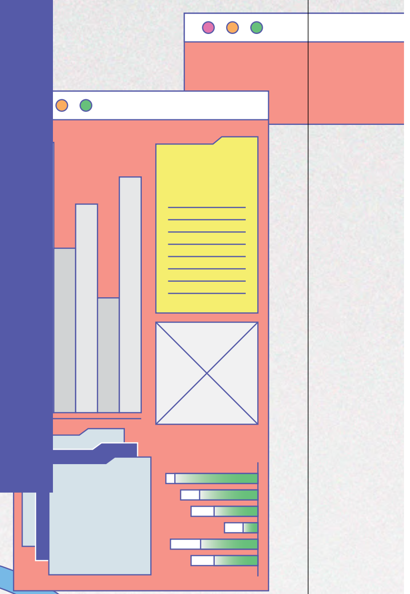

Illustrations by Ajun Chuah/Supercomputing Asia

SUPERCOMPUTING 21

After a hybrid event in 2020, the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC21)—the world's premier conference for all things supercomputing—is geared to return to an in-person event this 2021. To be held in St. Louis, Missouri, United States, the program proper will run from November 14–19 while exhibits will be open from November 15–18.

Bringing together researchers, industry experts and students, the conference promises an extensive schedule of plenary talks, panel discussions, poster sessions and networking opportunities. Students and early-career professionals especially can make the most out of the conference's job fair, mentor-protégé matching program and student competitions.

With the tagline 'Science & Beyond,' SC21 will delve into HPC's impact beyond the academic and research communities by spotlighting supercomputing breakthroughs in different sectors. There's something to pique everyone's curiosity, so book your spot to see how HPC makes a difference in everyday lives!

For more information, visit https://sc21.supercomputing.org/

WHERE

ST. LOUIS, MISSOURI, US

WHEN

NOVEMBER 14–19, 2021

HPC ASIA 2022

For its fifth edition, the International Conference on High Performance Computing in Asia-Pacific Region (HPC Asia) is organizing a hybrid conference in 2022, combining an in-person and virtual setup from January 12–14.

While the COVID-19 situation remains uncertain, conference organizers—the Information Processing Society of Japan, the RIKEN Center for Computational Science (R-CCS) and Kobe University—are hoping to host an onsite event at Kobe University's Integrated Research Center and neighboring R-CCS building. The latter is home to the world's fastest supercomputer, Fugaku, offering an enticing treat for attendees should the onsite event be possible.

First held in 2018, the upcoming conference remains dedicated to fostering the exchange of ideas as well as highlighting research advances and issues in the field of HPC. Through a combination of paper presentations and workshops, the much-awaited event will explore an interesting mix of algorithms and models; data storage and visualization; and memory and network architectures. Save the date and connect with leading HPC researchers in the region and globally!

For more information, visit http://www.hpcasia2021.org/

WHERE

KOBE, JAPAN AND ONLINE

WHEN

JANUARY 12–14, 2022

SUPERCOMPUTING ASIA 2022

Following the successful run of its first-ever digital conference in March 2021, SupercomputingAsia 2022 (SCA22) is gearing up for a big return through a hybrid event from 1–3 March in 2022. Since its inaugural edition in 2018, Asia's premier HPC conference has annually welcomed a diverse array of participants from over 20 countries across the globe. Attendees can anticipate the upcoming event to serve up exciting new insights, as HPC industry leaders and academic experts from both the public and private sectors convene to discuss the latest technological trends in supercomputing and its related topics.

Stay tuned for more details!

For more information on past conferences, visit https://sc-asia.org/

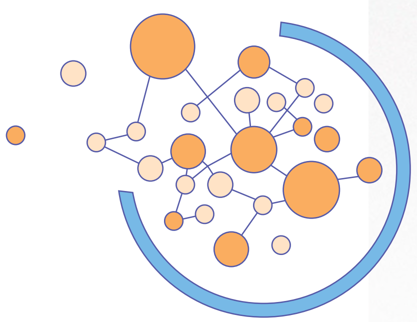

WHERE

SINGAPORE AND ONLINE

WHEN

MARCH 1–3, 2022

SOUTHEAST ASIA'S SUPERCOMPUTING SIGHTS

Take a tour of the region's HPC highlights

International or even regional travel might not be possible right now, but Southeast Asia still has much to offer the world when it comes to supercomputing.

By **Andre Encarnacion**

Illustrations by Lam Oi Keat/*Supercomputing Asia*

From the majestic stupas of Borobudur to the stunning beaches of Boracay, Southeast Asia boasts some of the most impressive tourist attractions in the world. Home to 655 million people speaking over 1,000 languages and dialects, the region is also one of the most diverse in the world—as well as one of the fastest growing.

With international tourism at a standstill thanks to COVID-19, Southeast Asia has been hit particularly hard. Despite the gloom, there have been bright spots, including increased regional cooperation and breakthroughs in COVID-19 research, many of which were enabled by supercomputing.

With the full opening of borders possibly a year or more away, join us on a virtual tour of some of the most outstanding supercomputing sights in Southeast Asia, from Thailand and Indonesia to Singapore and the Philippines.

Indonesia

Supercomputer's name	**BINUS Supercomputing Cluster**
Location	Bina Nusantara (BINUS) University AI R&D Center Jakarta, Indonesia
Fast facts	Developed to service BINUS University's genomics and deep learning needs, this workhorse is used in research projects tackling everything from cancer biomarkers to distance learning. With its versatility, the cluster is poised to propel BINUS University into a center for collaborative biomedical and artificial intelligence (AI)-focused research in all of Indonesia.
Top speed	25.78 teraFLOPS
Compute cores	34 Intel Xeon CPUs
Accelerators	11,712 NVIDIA Tesla and GeForce GPUs
Highlights	The BINUS AI R&D Center team recently used this cluster to process the first genome-wide association study for the Indonesian population, which aimed to discover the relationship between specific genetic variations and the country's most serious diseases. Although Indonesia has the fourth largest population in the world, with over 270 million people spanning 1,340 different ethnic groups, its people remain underrepresented in global genetics studies.

Thailand

Supercomputer's name	**TARA**
Location	National Science and Technology Development Agency (NSTDA) Supercomputer Centre Pathum Thani, Thailand
Fast facts	TARA is a vision of things to come for supercomputing in Thailand. Launched in 2019 to serve the NSTDA's R&D activities, it is the pilot cluster of the NSTDA Supercomputer Center (ThaiSC). In 2022, LANTA, a five petaFLOPs cluster, is expected to be in operation, expanding HPC service capability to the national scale.
Top speed	250 teraFLOPS
Compute cores	4,320 Intel Xeon CPUs
Accelerators	30 NVIDIA Tesla V100 GPUs
Highlights	Thailand was the first country outside of China to record cases of COVID-19 and TARA was used to analyze the initial strain of the SARS-CoV-2 virus. The information gleaned from the sequencing studies helped the government mount a swift and precise response in the pandemic's early days.

"These timely results led to a precise response to counter the spread of the original COVID-19 strain, possibly being one of the factors that helped Thailand avoid a surge of cases early on in the outbreak."

Dr Manaschai Kunaseth
Team Leader
NSTDA Supercomputer Center (ThaiSC)

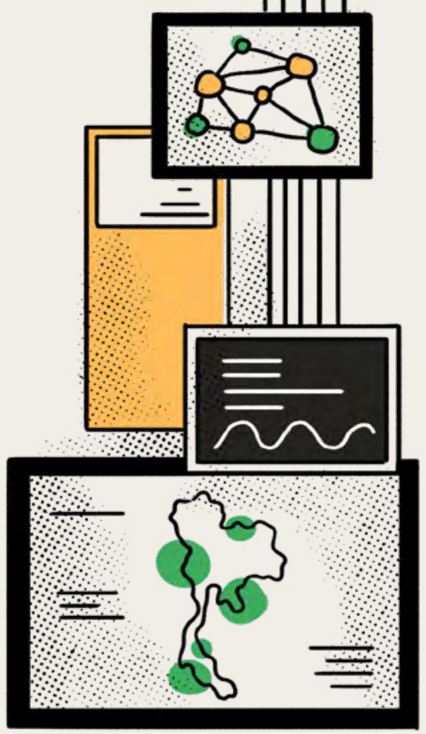

Singapore

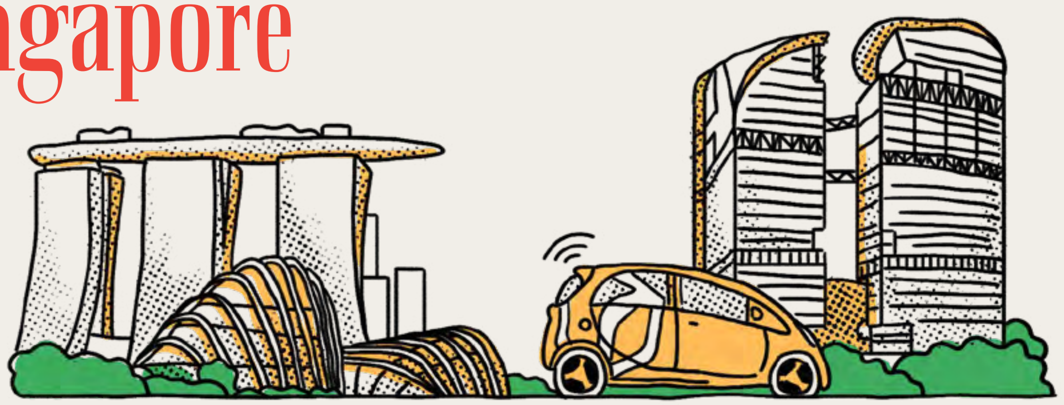

Supercomputer's name	**ASPIRE 1**
Location	National Supercomputing Centre Singapore
Fast facts	Launched in December of 2016, Singapore's Advanced Supercomputer for Petascale Innovation Research and Enterprise (ASPIRE 1) can be considered the elder statesman of Southeast Asia's supercomputing squad. After supporting a broad range of projects from self-driving cars to weather modeling, Singapore's National Supercomputing Centre (NSCC) announced in 2021 that they were investing S$40 million to build a 10 petaFLOPS successor to this legendary cluster.
Top speed	1 petaFLOPS
Compute cores	30,912 Intel Xeon CPUs
Accelerators	128 NVIDIA Tesla GPUs
Highlights	Seeking to overcome the limits of relying on subjective patient examinations to diagnose and manage psychiatric illnesses, researchers at the Singapore Bioimaging Consortium (SBIC), Agency for Science, Technology and Research (A*STAR), are using ASPIRE 1 to process psychometric, neuroimaging and other data to find quantitative biomarkers that better identify and help treat patients of diseases like schizophrenia.

The Philippines

Supercomputer's name	**Super Jojie**
Location	Asian Institute of Management Manila, the Philippines
Fast facts	Acer-developed Super Jojie is the beating heart of the Analytics, Computing and Complex Systems (ACCeSs@AIM) lab at the Asian Institute of Management. As the fastest of its kind in the country, it is currently being used to tackle national challenges like water shortages and educational performance during the COVID-19 pandemic. Currently optimized for AI applications, it is being upgraded to enable it to take on more Big Data workloads.
Top speed	1.2 petaFLOPS (single-precision)
Compute cores	48 Intel Xeon CPUs
Accelerators	96 NVIDIA GeForce GPUs
Highlights	Together with partners from government and the private sector, the ACCeSs@AIM team used Super Jojie to develop and deploy a contact tracing algorithm that prioritizes compromised individuals based on comorbidities, age and 'superspreader' potential—an appropriate use for a supercomputer that is named after an employee who secured the supercomputer's quick release from a customs checkpoint, known to her colleagues as a superwoman who can get difficult things done!

ACCELERATING HPC ADOPTION ACROSS ASEAN

The state of supercomputing in Southeast Asia

High performance computing could help the many talented scientists, engineers and businesses in Southeast Asia reach their full potential.

By **Tim Hornyak**

Illustrations by Shelly Liew/*Supercomputing Asia*

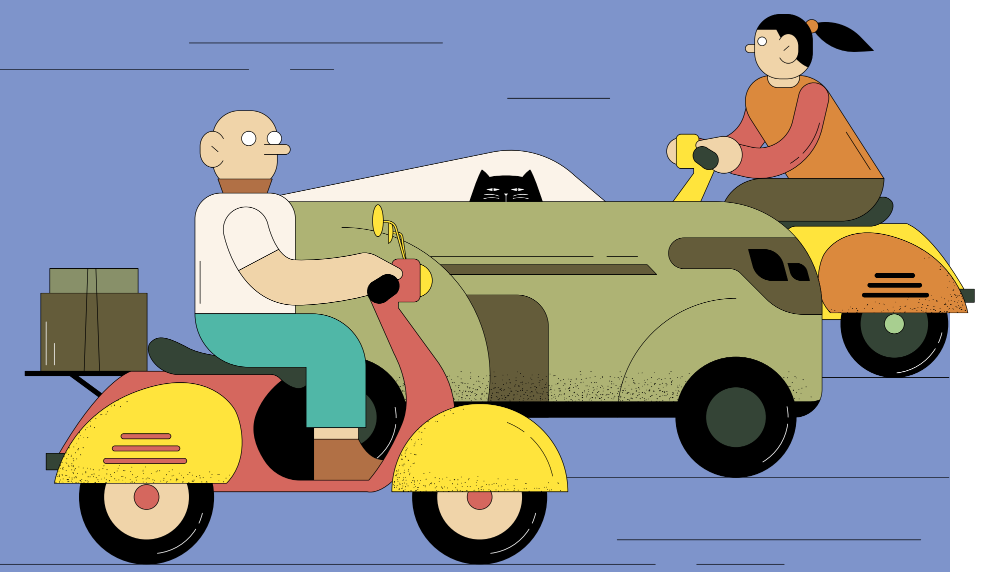

It goes without saying that your morning commute in Hanoi, Vietnam, would be vastly different from what it would be in Houston, Texas. There were over 276 million vehicles on US roads in 2019, a figure that works out to 842 cars per 1,000 people. In contrast, motorbikes dominate the roads in Vietnam, which has just 23 cars per 1,000 people—nearly 37 times less.

Yet when it comes to training autonomous vehicles (AV), the vast majority of data comes from the US and Europe. VinFast, a Vietnamese automotive startup with its sights set on the global market, is hoping to use supercomputers to bridge the difference. "When I came back to Vietnam, I had to relearn how to drive here—the traffic conditions are very different from the US," said Dr. Hung Bui, an artificial intelligence (AI) researcher formerly with Google DeepMind and Adobe Research, and currently the director of VinAI, the AI research arm of VinFast's parent company, Vingroup.

"After a while I got the hang of it, but it got me thinking a machine probably will do an even better job—Vietnam's driving conditions provide the ultimate challenge for systems trying to reach Level 5 autonomy," he said, referring to fully autonomous cars that will not even require steering wheels.

Helping VinFast along is the country's most powerful AI supercomputer: VinAI's NVIDIA DGX SuperPOD. Branded as the world's first cloud-native supercomputer, the DGX SuperPOD will be used to retrain VinAI's driving perception system as new data arrives every 24 hours, Hung said.

SUPERCHARGING R&D

World-leading AV systems built by private companies in Southeast Asia are just one of the many ways that high performance computing could impact the region. In light of the ongoing COVID-19 pandemic, for example, one of the most pressing needs is in drug development and healthcare. In Japan, the Fugaku supercomputer has been used to model the behavior of virus particles spreading among people.

"High performance computing provides researchers with tremendous power to address many grand scientific challenges," said Dr. Rossen Apostolov, executive director of BioExcel, a Centre of Excellence for Computational Biomolecular Research at the PDC Centre for High Performance Computing, KTH Royal Institute of Technology in Stockholm, Sweden. "Supercomputers allow screening of thousands of biomolecules in a matter of days for developing a cure or a vaccine against the virus. Similarly, modeling the propagation of a viral infection through a population allows for timely and effective measures to reducing its spread."

Using supercomputers to better understand the coronavirus comes amid growing HPC activity in the hemisphere. June 2021 will mark a year since Asia returned to the leading position in the TOP500 supercomputer rankings, with Fugaku—built by Fujitsu and research center RIKEN—bringing Japan back to No. 1 for the first time since 2012. Aside from US machines, the intervening decade has also seen Chinese systems—Sunway TaihuLight and Tianhe-2A—leading the pack, resulting in a greater Asian presence in the field of high performance computing.

The Association of Southeast Asian Nations (ASEAN) is represented in the ranking with Singapore's four HPC systems in the November 2020 list. While the TOP500 is only one way of taking the pulse of developments in HPC, Asia's growing role is becoming manifest in a number of national and international initiatives.

OVERCOMING SUSTAINABILITY CHALLENGES

In this flurry of activity, however, HPC developments in the ASEAN bloc have been lopsided, with Singapore taking up the lion's share of construction and investment. Dr. Piyawut Srichaikul, co-chair of the ASEAN HPC Taskforce and chief executive of the National Science and Technology Development Agency's Supercomputer Center (ThaiSC), shared that only about half the countries in the bloc are actively pursuing HPC goals.

"It's been some 30 years since HPC become a generic term, but I don't think ASEAN has an HPC regional landscape yet," Srichaikul told *Supercomputing Asia*. "I recall meeting colleagues in many gatherings related to HPC on both the infrastructure and application community sides from Singapore, Malaysia, Thailand, Philippines, Indonesia, Vietnam, etc. But, to my knowledge, only Singapore has steadily progressed during the past 30 years."

Srichaikul attributes this to HPC infrastructure sustainability problems among bloc countries that have fallen behind Singapore, including Thailand. After various planning permutations, ThaiSC was founded in 2019 with a mission of providing HPC services for research and development in Thailand. Located in Thailand Science Park north of Bangkok, the center has used its initial TARA system to service local users, take part in screening existing drugs for efficacy against COVID-19, and analyzing the coronavirus genome. ThaiSC is now procuring a new petaFLOPS flagship system, LANTA, that will be Thailand's first large-scale liquid-cooled supercomputer when launched in 2022, Srichaikul added.

"Competing to be on the [TOP500] list is not our prime objective," said Srichaikul. "This new system is small compared to those in the top rankings, but it does not matter; we look at HPC as a fabric of science and technology development.

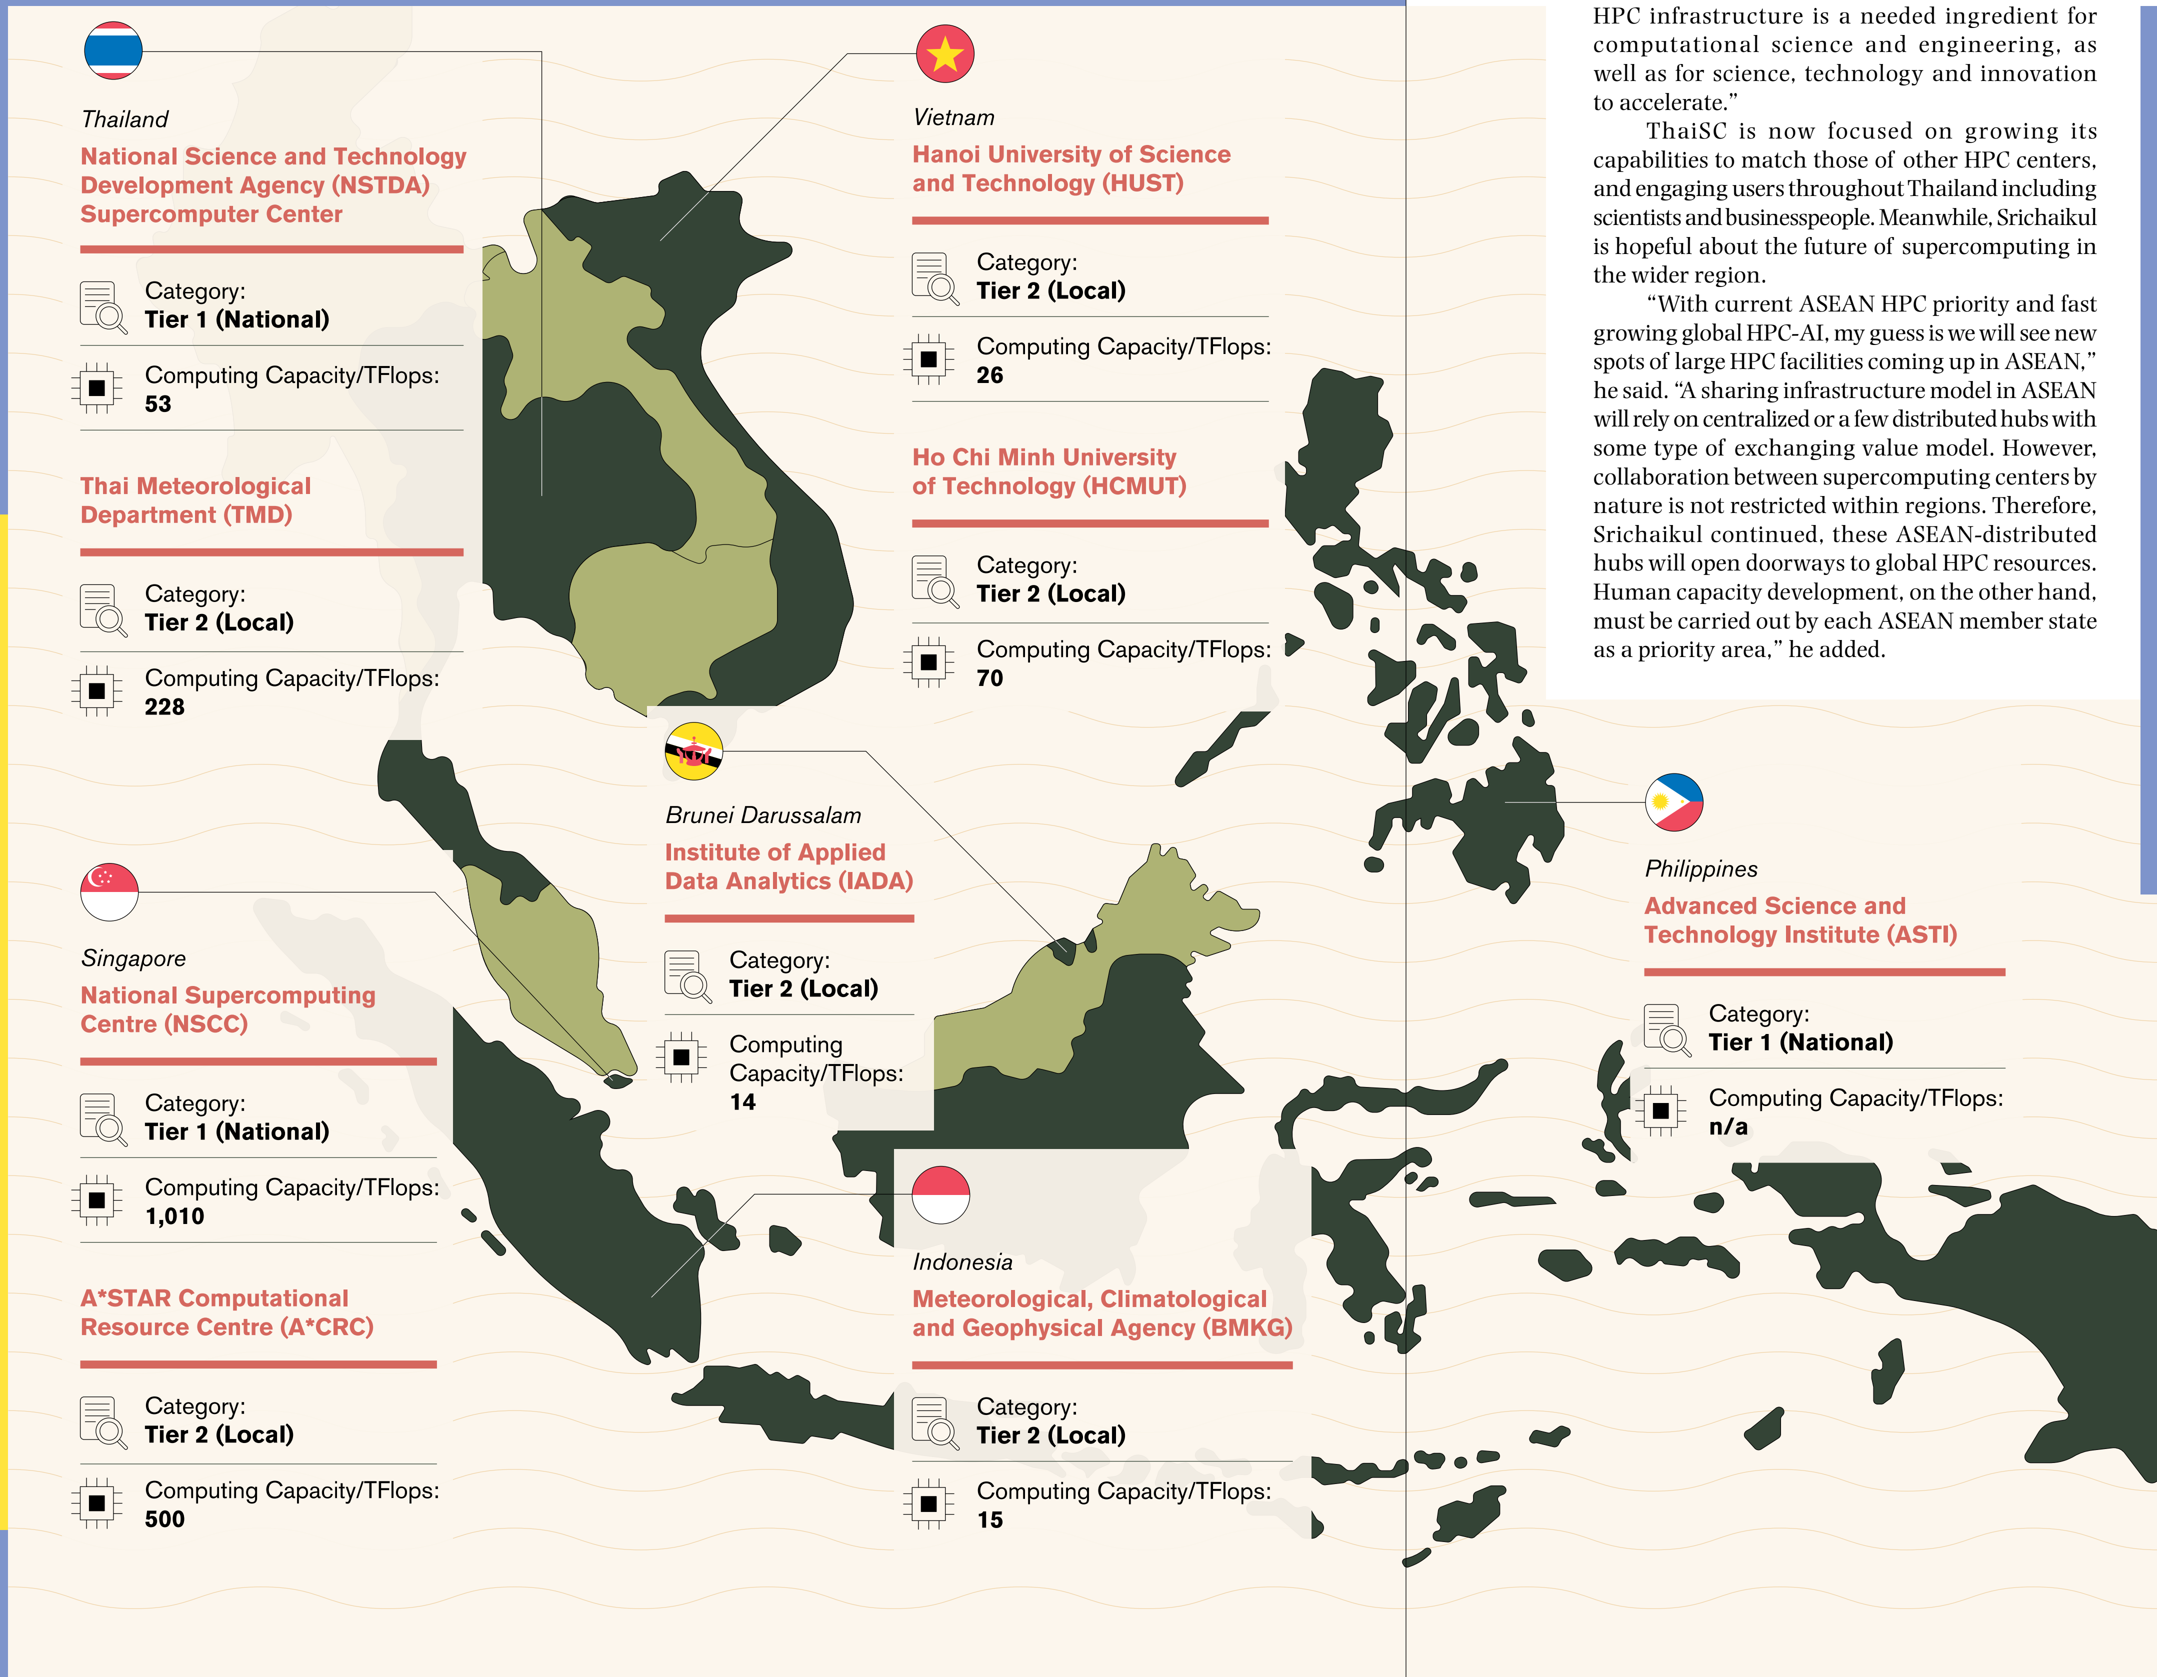

HPC infrastructure is a needed ingredient for computational science and engineering, as well as for science, technology and innovation to accelerate."

ThaiSC is now focused on growing its capabilities to match those of other HPC centers, and engaging users throughout Thailand including scientists and businesspeople. Meanwhile, Srichaikul is hopeful about the future of supercomputing in the wider region.

"With current ASEAN HPC priority and fast growing global HPC-AI, my guess is we will see new spots of large HPC facilities coming up in ASEAN," he said. "A sharing infrastructure model in ASEAN will rely on centralized or a few distributed hubs with some type of exchanging value model. However, collaboration between supercomputing centers by nature is not restricted within regions. Therefore, Srichaikul continued, these ASEAN-distributed hubs will open doorways to global HPC resources. Human capacity development, on the other hand, must be carried out by each ASEAN member state as a priority area," he added.

PARTNERING FOR GROWTH

"The ASEAN region's interest in HPC is growing significantly and a number of ASEAN member states are attempting to increase super-compute capabilities and capacities to meet the ever increasing demand," said Mr. Tay Kheng Tiong, chief executive officer of the Computational Resource Centre (A*CRC) at Singapore's Agency for Science, Technology and Research (A*STAR). "However, due to various constraints, notably economics and other more pressing priorities of the individual nations, the urgency is not quite there."

"On the other hand, Singapore, Thailand, Malaysia and Vietnam have taken strategic initiatives at the national level to provide the necessary technology and solutions to meet individual countries' key research domains that require huge computational resources," Tay said.

A significant step forward for ASEAN's HPC growth has been increasing international collaboration, namely with Japan and Europe. For example, the Singapore National Supercomputing Centre (NSCC) has tied up with RIKEN's Center for Computational Science (R-CCS), which manages Fugaku. These agreements will give scientists in the ASEAN region access to world-class computing

> We want to help them accelerate and quickly reach the level of supercomputing they deserve without going through all the pains that we experienced.
>
> Professor Satoshi Matsuoka
> Director of RIKEN's Center for Computational Science (R-CCS)

infrastructure without the enormous investments required to build leading computing tools.

"Countries like Thailand can make significant advances if scientists can get their hands on these advanced infrastructures, but it takes a lot of experience and backing and people skills to drive them—in Japan, we've been doing this for 50 years," said Professor Satoshi Matsuoka, director of R-CCS. "We want to help them accelerate and quickly reach the level of supercomputing they deserve without going through all the pains that we experienced. That way, their infrastructure can be on par with the science and technology skills their engineers have."

There are several ways scientists in ASEAN can avail themselves of resources in Japan. R-CCS is working with Japan's Research Organization for Information Science and Technology (RIST), which allocates HPC cycle time in Japan on a competitive basis, to create a new framework specific to foreign research projects, including those from ASEAN, according to Matsuoka. ASEAN researchers can also collaborate directly with R-CCS as research partners, a process that requires a memorandum of understanding but no competitive application for cycle time. Meanwhile, R-CCS has been working with counterparts in Singapore to make using Fugaku as easy as possible with the made-in-Japan Grid Data Farm (Gfarm) data file system. It lets users run large volumes of data in HPC applications without having to transfer them from one machine to another.

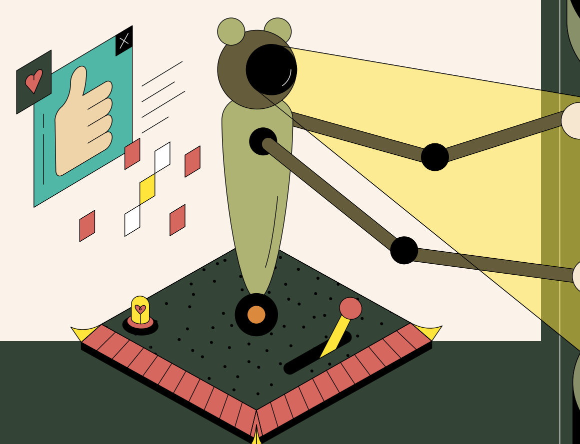

The near-term objective is to establish the inaugural ASEAN shared facility on Fugaku.

Mr. Tay Kheng Tiong
Chief executive officer of the Computational Resource Centre (A*CRC)

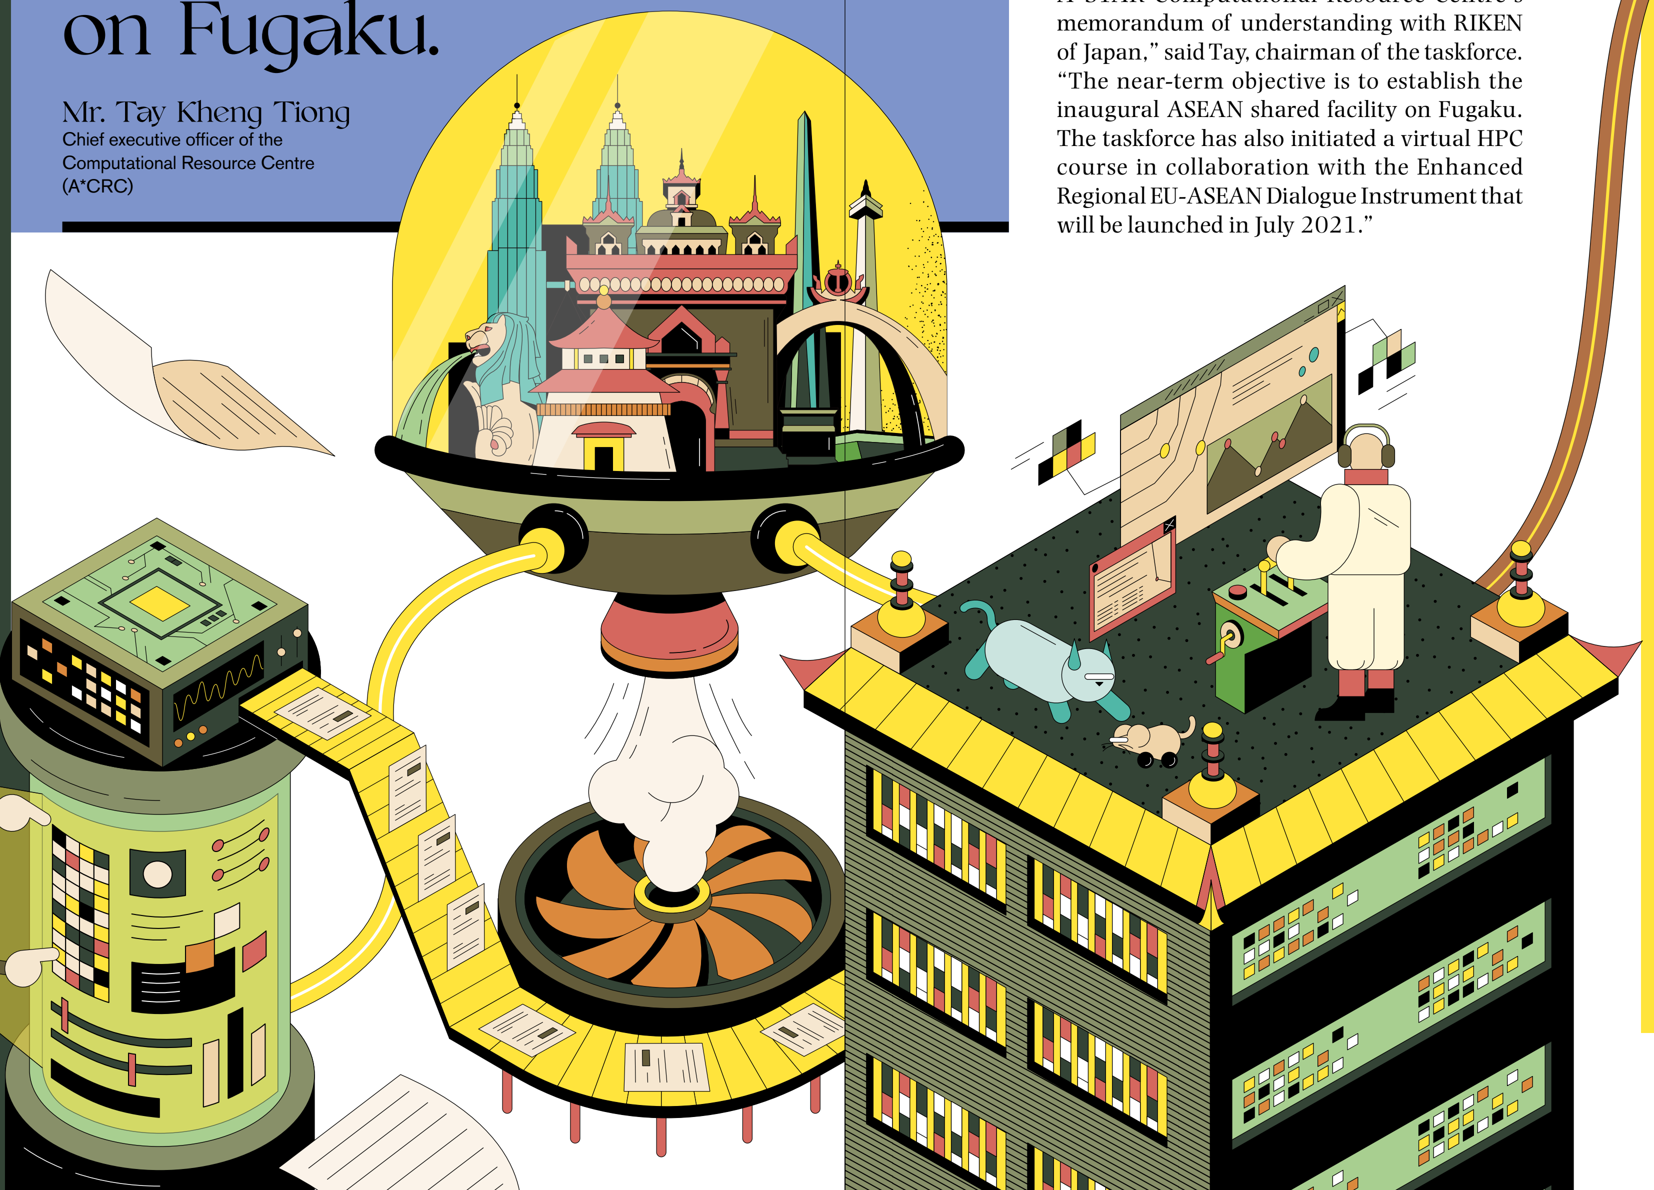

BETTER TOGETHER

A key driver of this collaboration is the ASEAN HPC Taskforce. It was formed in early 2018 to establish a framework of HPC collaboration among ASEAN countries to identify common areas of research. One of the Taskforce's key recommendations was a shared HPC facility, a massive undertaking that will require not only international collaboration, but investments and human resources.

"So far, the taskforce has managed to work with the RIKEN HPC experts to set up a testbed of shared infrastructure on Fugaku through A*STAR Computational Resource Centre's memorandum of understanding with RIKEN of Japan," said Tay, chairman of the taskforce. "The near-term objective is to establish the inaugural ASEAN shared facility on Fugaku. The taskforce has also initiated a virtual HPC course in collaboration with the Enhanced Regional EU-ASEAN Dialogue Instrument that will be launched in July 2021."

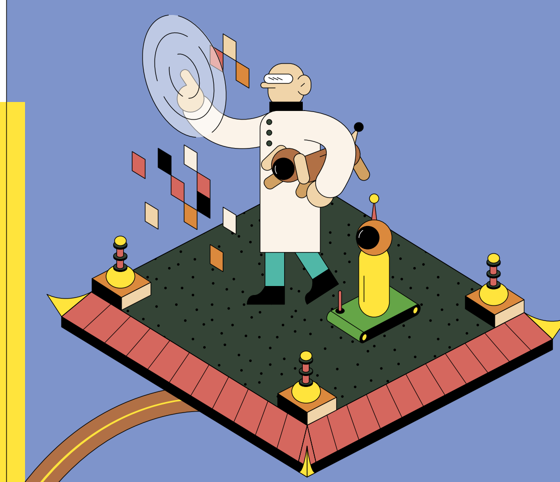

That course is the EU-ASEAN High-Performance Computing Virtual School 2021. It's a five-day virtual seminar in July 2021 in which HPC experts will instruct 60 participants from the ASEAN region in the basics of HPC design and applications in areas such as climate science and COVID-19. Students will be ranked on their performance in practical exercises, and the top ten will receive a digital certificate of excellence and be guaranteed a place at the next EU-ASEAN HPC school in Thailand.

"Grand scientific challenges require concerted efforts beyond the capabilities of a single research group," added KTH's Apostolov, one of the school's coordinators. "It is important that software tools, expertise and resources are shared and exploited effectively in a collaborative manner by the global communities. The EU-ASEAN HPC initiative is a necessary and great step towards knowledge exchange and technology uptake between the regions."

With growing international collaboration, stronger government backing and greater international recognition of HPC as an engine for national science and technology programs, it's clear that there's a bright future for supercomputing in the bloc.

"There are a lot of things in motion," said Matsuoka, referring to transborder initiatives such as the virtual school, where he will be a keynote speaker. "If ASEAN countries were individual nations, it would be very hard to play the catch-up game. Because they're doing it together, I think they can advance much quicker."

THE SILVER LINING ON COVID'S DARK CLOUD

The role of HPC collaborations in the new normal

From cross-border collaborations to advancing accessibility, supercomputers play an important role in helping the world recover from COVID-19.

By **Sheryl Lee**

As COVID-19 spread across the globe in early 2020, flights halted, borders closed and people stayed home. But while the pandemic made national borders painfully apparent, it brought countries together, too. Researchers across the world sequenced and shared the SARS-CoV-2 genome; nations sent masks, testing kits and even their own doctors to the hardest-hit regions; and millions of dollars flowed in the name of international aid.

Likewise, the pandemic has drawn the high performance computing (HPC) community closer together over the past year, as people rose to the occasion and channeled their supercomputing powers for good. Supercomputers not only sequenced millions of viral genomes, but also helped us model how the disease is spread and what measures could possibly counteract it.

COMPUTING AGAINST COVID-19

Coming online several months early to begin working on such problems was Japan's Fugaku supercomputer, which was originally scheduled to begin operating only in 2021, shared Professor Satoshi Matsuoka, director at the RIKEN Center for Computational Science in Japan. Matsuoka was speaking at the inaugural HPC Center Leaders Forum, an expert panel at the SupercomputingAsia 2021 conference held virtually from March 2–4, 2021.

From screening existing therapeutics for repurposing into COVID-19 drugs, to detailed simulations of droplet aerosols modeling transmission, Fugaku was immediately put through its paces and harnessed to combat COVID-19, Matsuoka said. "You go from very microscopic atomic levels all the way to societal levels in various simulations in trying to provide solutions to the pandemic."

Building on this initial success, there are plans to open up access to Fugaku for international research partnerships, Matsuoka continued. "We have already developed a system that allows for very effective sharing of data across HPC supercomputers. We are testing this to see if we can expand it to Australia, Singapore and other Asian countries," he said. "With this system, [researchers from other countries] will be able to use Fugaku through the cloud. A file can be local to Singapore, with gigabytes of data transferred transparently in an instant."

ACCESS AND AWARENESS

In this spirit of openness, two unlikely partners have also ways to collaborate on HPC. Separated by more than 9,000 kilometers, the tropical island of Singapore and the Scandinavian country of Finland have nonetheless signed an memorandum of understanding to explore a high-speed fiber optic link between the two countries, as well as more secure ways of protecting data transfer through quantum technology, said National Supercomputing Centre Singapore's chief executive Associate Professor Tan Tin Wee.

"We have also focused a lot more on the societal impact of HPC and the translation aspects into applications, particularly for our research and industry communities," said Tan. He added that the ultimate goal is to reach out to as many researchers as possible, including students, and regardless of the amount of HPC needed, so that the access to HPC can be truly democratized.

Concurring with Tan, Dr. Piyawut Srichaikul, chief executive at Thailand's National Science and Technology Development Agency Supercomputer Center, further emphasized the importance of HPC literacy.

"[In Thailand], public awareness and understanding of HPC—from laymen to authority figures and politicians—is very little. [Lack of] human capital is a major problem that we've been facing for years now," he said. "As a result, it is difficult to determine HPC's direct impact on our economy and society, because ultimately it is the user who creates the impact." This conundrum has made it more difficult for the country to understand the benefits and importance of HPC as a whole, and invest in it, Srichaikul shared.

We have also focused a lot more on the societal impact of HPC and the translation aspects into applications, particularly for our research and industry communities.

Associate Professor Tan Tin Wee
Chief Executive
National Supercomputing Centre Singapore

LOVE YOUR SUPERCOMPUTING NEIGHBOR

While the public might not be aware of how supercomputers are contributing behind the scenes, one thing that has captured public imagination is the need to be environmentally friendly. After all, for supercomputers to truly be a force for good, they must not only benefit people but the environment as well, pointed out Mr. Kimmo Koski, CEO of the Finnish IT Center for Science. He explained how, in deciding on a location for the pre-exascale supercomputer LUMI, the consortium had zeroed in on Kajaani, a part of Finland with an abundant supply of hydroelectric power.

The availability of renewable energy aside, Kajaani's naturally chilly climate also makes free cooling possible all year round, Koski shared. "We also want to use the excess heat from the supercomputer to warm up the households in the area," he added. "This will make it very environmentally friendly and the carbon footprint will be practically [negligible]."

Wrapping up their discussion, the four leaders concluded that HPC will likely continue to permeate across all sectors of society and across borders. Accessibility and collaboration will be crucial going forward, coupled with an emphasis on sustainability, public awareness and support. The pandemic may have brought supercomputers to the forefront, but if they are to stay, they must be well-integrated into society as well.

Illustrations by Alexandra Valino, Lam Oi Keat/*Supercomputing Asia*

Setting a course for supercomputing

The many roads leading to HPC

From molecular modeling to formulating facial wash, the world of high performance computing has much to offer students from different fields and passions.

By **Jill Arul**

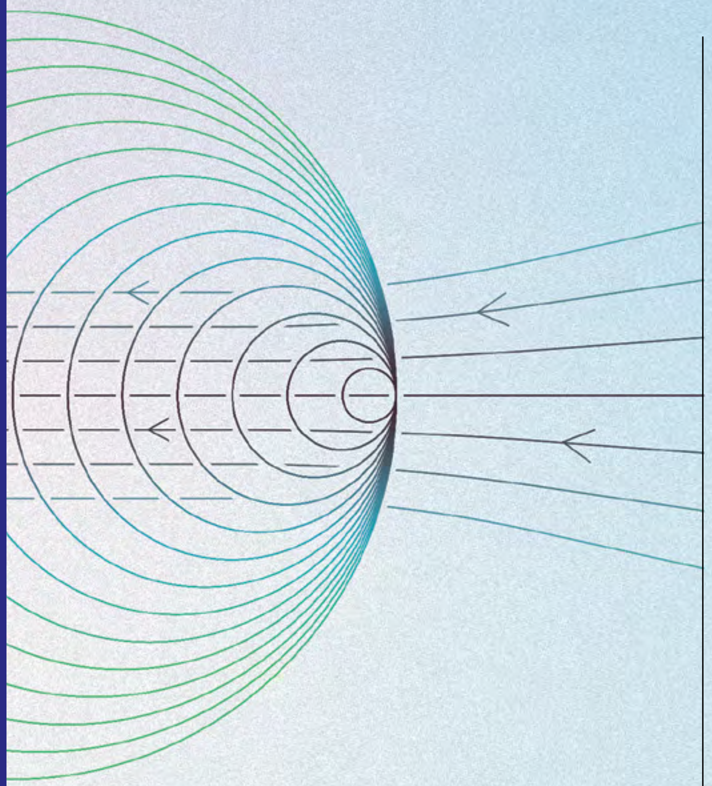

In 2020, as part of an international COVID-19 high performance computing (HPC) consortium, researchers in Korea analyzed 19,168 small molecules to identify eight potential drug candidates that could have therapeutic potential against COVID-19. On the humanities front, as art is digitalized and online data collected, researchers scale new heights with access to more data and computing tools than ever before.

From biomedical science to fine art, the impact of supercomputing is widespread and shows no sign of slowing down. With the global supercomputing industry on track to reaching US$49.4 billion by 2025, the future is bright for aspiring computational researchers who can expect to enter a thriving space.

Covering everything from life as a computational chemist to HPC for beginners, specialists and scientists at the education session of SupercomputingAsia 2021 on March 4, 2021 offered pearls of wisdom to students to prepare them for the many roads leading to HPC.

GOING DOWN A CHEMICAL PATH

For students looking to dive into the depths of HPC as computational researchers, there are several paths to choose from—a particularly popular route to take is that of a computational chemist.

Commonly used in chemistry applications, HPC allows scientists to compute models and study the structure and dynamics of various molecules. Demonstrating the far-reaching applications of molecular modeling, multiple institutes and HPC centers across the globe, including several right here in Asia, have come together to model SARS-CoV-2's structure and focus their computing power toward drug and vaccine development and testing.

Computational chemists start by building models that represent specific systems. Real-life laws and statistics are then applied to the modeled electrons, atoms and molecules in the system. Finally, a simulation is run to test potential theories.

"We move around from theory, to simulation and to experiment," Dr. Adrian Mak from the Institute of High Performance Computing (IHPC) at the Agency for Science, Technology and Research (A*STAR), explained. "The theory provides the equations for your simulation and your simulations provide interpretation for your experimental results."

With the potential to solve some of the world's most pressing issues, molecular modeling and computational chemistry are powerful aspects of HPC poised to address a new generation of challenges.

A SOLUTION FOR EVERY PROBLEM

Aside from experimental research, HPC can also potentially transform commercial products we use every day. When it comes to skincare and haircare, companies in the fast-moving consumer goods sector invest heavily in research to create the best products. L'Oréal, for example, is a particularly big spender, dedicating roughly US$1 billion to research and innovation in 2018.

For scientists stepping into the world of consumer care, exploring new computational methods and applying their research to commercial solutions can be a fresh experience.

"From looking at bond breaking and bond forming, I moved in the direction of studying physical interactions in soft materials like hair and skin, to design gentle and effective formulas for conditioners and facial cleansers," shared Dr. Freda Lim, currently a senior scientist and innovation lead at A*STAR's IHPC.

Alternatively, instead of focusing on one industry, a single specialization can offer solutions to multiple industries and issues. One example of this is computational fluid dynamics (CFD), where researchers study the flow of liquids or gases around an object or area. When applied to aerodynamics, environmental engineering and more, an understanding of CFD could open up a world of opportunity for students with a variety of interests.

From optimizing service reservoirs to mapping blood flow in stroke patients, Dr. Kang Chang Wei, deputy department director at A*STAR's IHPC, and his team harness CFD to provide sustainable solutions to issues that span many industries. Putting their expertise towards managing the pandemic, Kang and his team have even modeled how droplets propagate when people sneeze to potentially track the spread of the virus.

SUPPORTING ALL SECTORS OF SCIENCE

Across disciplines, researchers today have no shortage of data to collect as they pursue their projects. Medical researchers, for example, have access to a wealth of genomic information, with over 3,000 individual genomic resources, tools and databases available publicly on the internet alone. Meanwhile, ecologists have sensor networks in place to continuously collect millions of unique field observations like temperature and movement.

But even after obtaining such huge amounts of data, the work is far from done. Researchers need to be able to analyze the raw information and draw relevant, accurate conclusions. Enter HPC, a powerful multi-faceted tool that plays an important role in making sense of data.

> "DON'T BE AFRAID TO BREAK STUFF, DON'T BE AFRAID TO BREAK CODE—IN FACT, BREAK IT AND TRY TO REASSEMBLE IT AGAIN."
>
> DR. JERNEJ ZIDAR
> SENIOR HPC ANALYST, NATIONAL SUPERCOMPUTING CENTRE SINGAPORE

While some researchers are comfortable with HPC, others are not and may even actively avoid it—proving detrimental to their work. According to Ms. Julie Faure-Lacroix, science liaison agent at University Laval Center of Calcul Quebec in Canada, typically, researchers who regularly make use of HPC do so for studies that involve intricate processes like molecular modeling or aerodynamics.

However, as big data and artificial intelligence enter the humanities, HPC centers must be able to offer tools and guidance to a new demographic of users.

"[New users] don't want to learn what a terminal is or how to code properly," Faure-Lacroix explained. "They want tools that will give them access to supercomputers without the learning curve of suddenly having to become a computer scientist after years of being an artist."

TAKING THE FIRST STEP

As with most other career paths, technical mastery is needed for students to become well-equipped for a life in HPC. In particular, students should be comfortable with mathematics, have sufficient knowledge of at least one programming language and be familiar with major operating systems like Linux.

That said, attaining the necessary skills does not have to be difficult with affordable or free online courses readily available for students to get familiar with the basics of computing.

To get to know the ins and outs of the HPC community, students are encouraged to take on internships for a hands-on experience in a HPC center. Institutes also stand to gain from introducing bright-eyed and eager young people into their team.

"We will benefit from having fresh blood and the smart brains of energetic students who want to explore more and challenge us with different questions," said Dr. Wang Jingbo, a senior staff scientist at the National Computational Infrastructure in Canberra, Australia.

To truly seize all opportunities available to them, Dr. Jernej Zidar, a senior HPC analyst at the National Supercomputing Centre Singapore recommends that students enter HPC unafraid. Students are encouraged to jump into internships, hackathons, online courses and any opportunity to code.

"Don't be afraid to break stuff, don't be afraid to break code—in fact, break it and try to reassemble it again," he urged. "It gets frustrating now and then, but you learn so much in the process."

With so many possibilities ahead, the journey to a career in supercomputing can seem long and arduous—but there is help along the way.

POSITS: COMING SOON TO HARDWARE NEAR YOU

With applications in artificial intelligence and extensive hardware compatibility, high-accuracy posit arithmetic is set to rewrite the standard for computing in the years to come.

By **Erinne Ong**

Compared to humans, computers are faster and more accurate and precise—or at least, they are supposed to be. But dig a little deeper beneath the surface and you'll find inexactness and imperfection. Most of us who use them every day might not know this, but mathematical calculations on computers have a dirty little secret: they can't count.

Don't believe me? Try adding 0.1 and 0.2 using the standard format of floating point representation. While a human mind would quickly arrive at the answer of 0.3, a computer would end up with an answer like 0.2999998 or 0.3000002—close, but no cigar. The reason is that while humans understand that adding integers like 1 and 2 will make exactly 3, computers don't think in decimals and need to convert them to a format they do understand: binary floating point. This representation isn't perfect, leading to rounding errors and poor reproducibility.

If that seems like a let down, look no further because a new data type called posits now offer greater accuracy for less computational resources. Pioneered by Dr. John Gustafson, a professor at the National University of Singapore, posit arithmetic and its applications have advanced in leaps and bounds since their introduction in 2017. Researchers in Asia and globally have kept the ball rolling, ushering in new developments that may soon enable posits to light up our computers in full force.

A PRIMER ON POSITS

There is a finite way of storing numbers on a computer. As bits can only be in zero or one states, converting the decimal numerical system of zero to nine means needing more digits—more bits—to represent the same number. That is computationally taxing and highly energy-inefficient for non-integer real numbers like fractions and square roots.

As the dominant standard accepted by the Institute of Electrical and Electronics Engineers (IEEE), the floating-point format allocates bits to the significand and exponent components in a predetermined way. In a trade-off between accuracy and range, approximations are inevitable as computers run out of bit-width and calculations become littered with unusable elements like negative infinity or not-a-number representations.

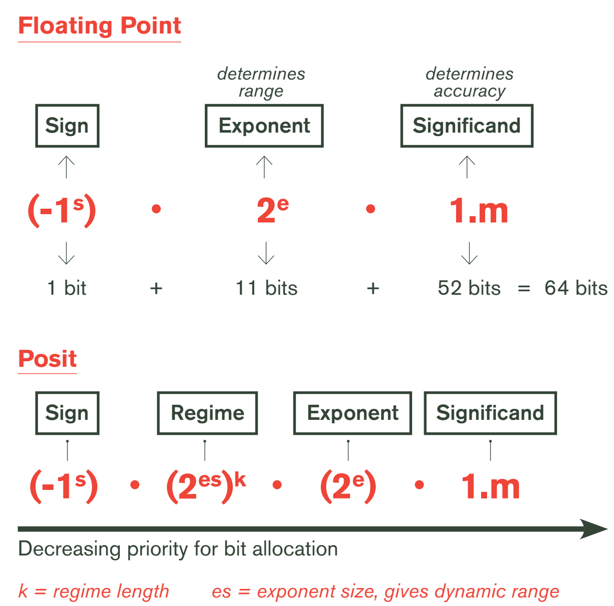

While floating points have fixed bit allocations for each part of the expression, the posit format allows bits to be distributed in a flexible way, with priority given to the sign and regime terms. This leads to posits' characteristic dynamic range and high accuracy, overcoming the limitations of floats that have fewer bits available for the significand. A standard 64-bit posit has a 60-bit significand for most computations, compared to only a 53-bit significand for IEEE standard floats, making the posit over 100 times as accurate while using the same number of bits.

"It assumes that you want the same accuracy, no matter how big or small the number is. In reality, you don't really want that—you want more accuracy for the commonly used numbers that are close to one," Gustafson said.

This insight underlies the introduction of a regime bit in posit representation. Unlike floats, the number of bits per part is not fixed, with priority given to the regime term ahead of the exponent and fractional components.

> "We might be able to build hardware that's actually not just a better design for accuracy, but much faster than anything we have right now on supercomputers."

Professor John Gustafson
National University of Singapore

"The regime bit allows the accuracy to taper automatically. This sliding scale works very elegantly to create the effect of putting more emphasis where you have the common numbers," explained Gustafson.

In effect, this enables choosing the right balance between dynamic range and accuracy to suit the task at hand. Because the regime field only requires two bits for most computations, posits are well-optimized—assuring that the bits are allocated where they are needed most.

BETTER ANSWERS, FEWER BITS

With posits, complexity is often turned into something much simpler. For computers, multiplication is a particularly taxing operation and prone to inaccuracies, especially when involving non-integer real numbers.

Such is the case with the logarithmic number system, for example. While converting to logarithm allows for addition instead of multiplication, these operations have always relied on approximations instead of the exact value. The result is a bunch of long decimals that have to be cut off to fit within a computer's bit limits, often leading to rounding errors.

But in a recent breakthrough for posit arithmetic, a multiplication table of real numbers can be matched up and turned into an addition table—using small integers at that. By foregoing decoding steps or mapping back to the original representation, calculations are made simple and low power yet provide exact answers. Such perfect mapping works best at the low precisions used by artificial intelligence (AI) applications, shared Gustafson.

While it may sound like a negative attribute, low precision refers to systems running on 16 or fewer bits, compared to single-precision designated at 32 bits. In other words, posits are doing much more with less.

Consider the balancing act of compression, such as for an image or audio file, where file size is reduced to as small as possible without evidently sacrificing image sharpness or sound quality. Similarly, 16-bit posits are essentially lossless when they reproduce data signals, whereas the deeply entrenched IEEE 32-bit floats lose too much information.

By providing better answers with fewer bits, Gustafson highlighted that posits can spark an immense jump in computational power and speed for AI applications. "We might be able to build hardware that's actually not just a better design for accuracy, but much faster than anything we have right now on supercomputers," he added.

One particular technology that would stand to benefit most are deep neural networks (DNN), multi-layered algorithms that attempt to emulate how the human brain processes information. While simpler AI models analyze data when given a certain ruleset, neural networks learn on their own by manipulating and making sense of tons of training data, assigning labels the same way we distinguish between different types of dogs or how doctors identify damaged tissues in medical scans.

As such, DNN applications are typically developed for highly complex tasks like image processing and autonomous driving. However, Gustafson noted that these systems are extremely resource-intensive. "The training can go on for days, just for the simplest task, so you're talking about many kilowatt hours of energy consumption," he explained. "It's a bottleneck."

Posits have the potential to markedly enhance efficiency in DNN training and inference—the latter referring to the predictive capabilities of AI models, which put their learning to work in finding patterns in new data.

Working with Gustafson, a US research team showed that 8-bit posits outperformed equivalent floats and had comparable results with 32-bit floats in DNN inference. Classification accuracy ranged from a promising 86 to 99 percent for various tasks, such as distinguishing malignant versus benign breast cancer tumors and identifying numbers of varying handwriting styles.

Aside from high accuracy, posit arithmetic closely matches the representation of a certain function that is used often in neural networks, accelerating learning and computing with less resources.

"The ability to use fewer bits means that you're going to use a lot less energy and a lot less space on a chip, and all the costs go down," Gustafson said.

SMOOTHING OVER THE SWITCH

Reaching the next frontier in high-accuracy and energy-efficient computing will depend not just on posit arithmetic but the underlying platforms that can support it. According to Gustafson, researchers around the world have already built math libraries that will support conversion routines between one data type to another, like floats to posits and vice versa.

At its current stage, posit representation functions as if it were the more widely used floats, but provides much improved calculations, even for complex operations such as exponents and trigonometric functions like sines. As part of the posit standard, which Gustafson is developing and soon set to publish officially, such conversion would pave the way for wider uptake of the posit format.

"Having a standard you can point to and having portability between all the different systems is key to having at least one major company embrace posits to the point where it's supported on what they produce as a product," he said.

Among the recent advancements in this space is PACoGen, created by University of Hong Kong researchers for generating posit arithmetic cores that can perform basic operations. As these cores form part of the circuitry of computers, PACoGen serves as a pipeline for sketching out the integration of posit processing units (PPUs) into hardware like microchips, showing how posit formats can fit within various technical specifications.

Another crucial piece of the puzzle is posits' compliance with RISC-V instruction set architecture, an open-source set of design specifications that ensure compatibility across different processors.

"The breakthrough was realizing I could give up on divides as a perfect operation. If I can make add, subtract and multiply perfect, then I don't have to have divides be perfect because they're less common," explained Gustafson. Like picking only the battles that can be won, this optimization led to a clean number system that fit in a neat size and presented the same ordering as integers.

While posits were conceptualized as drop-in replacements for floats, Gustafson acknowledged that there will be a transition period where users can switch between the two formats, akin to having different processing modes.

From the frontiers of computing to mainstream tech, this transitional phase has already been set in motion. At ThoughtWorks Technologies India, for example, scientists replaced floats with PPUs on their Rocket Chip. Evaluated on the RISC-V test suite, they found that posits minimized the amount of unusable representations appearing in the calculations. At the same time, the PPUs still supported RISC-V floating point extensions—additional codes allowing certain functions to be carried out as needed.

As new processors are already being designed, Gustafson envisions that posits would become the dominant format supported by the hardware, while floats become more of a rarely used legacy tool. Floats would still be available, but only through software extensions and libraries. For him, overcoming intellectual inertia and making the switch to posits could happen as soon as the next decade.

"I think posits will gradually be adopted and they will compete, and therefore computing will move forward in the same way," concluded Gustafson. "That kind of improvement is gradual and I don't think it has to be that painful. I hope posits will eventually take over the world."

Business Bytes

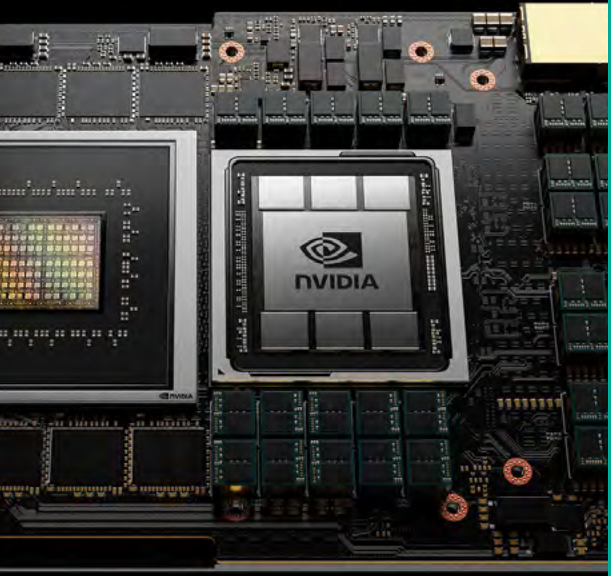

Credit: NVIDIA

GRACE CPU ADDED TO NVIDIA'S STABLE OF CHIPS

Chipmaker NVIDIA Corporation has unveiled its first-ever data center central processing unit (CPU). Named Grace, the chip set combines CPU cores from chip design firm Arm Ltd. and NVIDIA's graphical processing units (GPUs).

Grace-based systems are touted to offer up to tenfold faster processing than current state-of-the-art systems running on x86 CPUs. To achieve greater bandwidth than existing servers, NVIDIA also incorporated better memory architecture and CPU-GPU connections of up to 900 gigabytes per second.

This highly specialized processor is geared to push the frontiers of HPC and AI, according to NVIDIA CEO and founder Jensen Huang. By analyzing copious amounts of data, Grace is set to keep pace with the billions of parameters that leading AI models now contain.

"NVIDIA has designed Grace as a CPU specifically for giant-scale AI and HPC," Huang added. "Coupled with the GPU and data processing unit (DPU), Grace gives us the third foundational technology for computing and the ability to re-architect the data center to advance AI. NVIDIA is now a three-chip company."

AI-CAPABLE SUPERCOMPUTING REACHES NEW HEIGHTS

The Swiss National Supercomputing Centre (CSCS), multinational IT company Hewlett Packard Enterprise (HPE) and chipmaker NVIDIA Corporation are creating a world-leading supercomputer capable of performing AI tasks. Named after the highest mountain range in Europe, the Alps system will similarly be making its name in the supercomputer history books when it comes online in 2023.

The system is set to use an underlying HPE Cray EX supercomputer architecture for analyzing complex data sets, which will be further enhanced by NVIDIA's HGX supercomputing platform of high-speed, high-bandwidth CPUs and GPUs.

Should Alps live up to the hype, its technical capabilities are expected to have a wide range of applications in cutting-edge research. This could lead to better climate models and forecasting systems, comprehensive molecular models for drug discovery, and faster natural language processing (NLP) tools, among other use cases. The developers claim that the system's processing power could train the world's largest NLP model in just two days, seven times faster than the current leading AI-capable supercomputer.

Envisioned to be a general-purpose system, Alps will be made available not only for CSCS users in Switzerland but also for researchers around the globe.

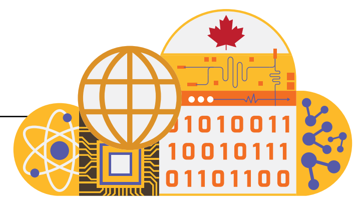

QUANTUM CLOUD SERVICE LEAPS ONTO SINGAPORE'S SHORES

Leap, a quantum cloud service provided by Canadian company D-Wave Systems Inc., has made its way to Singapore. Now launched in 38 countries globally, Leap brings a host of D-Wave computing solutions, geared towards accelerating digitalization and integration of hybrid quantum applications across various industries.

Available for real-time access are the company's Advantage quantum computer, the Quantum Application Environment, as well as developer resources such as learning tools, sample codes and an integrated development environment. The Advantage system provides significant processing power for tackling large datasets and complex problems in an easy-to-navigate platform.

Over 250 early-stage applications have been built with D-Wave computing systems, with the potential to add value in sectors such as logistics, manufacturing, finance and pharmaceuticals.

"We look forward to working with businesses in Singapore to reap the performance advantages, efficiencies, speed and scale that quantum computing provides to enable the country's industries to remain competitive in the region and globally," said D-Wave CEO Alan Baratz.

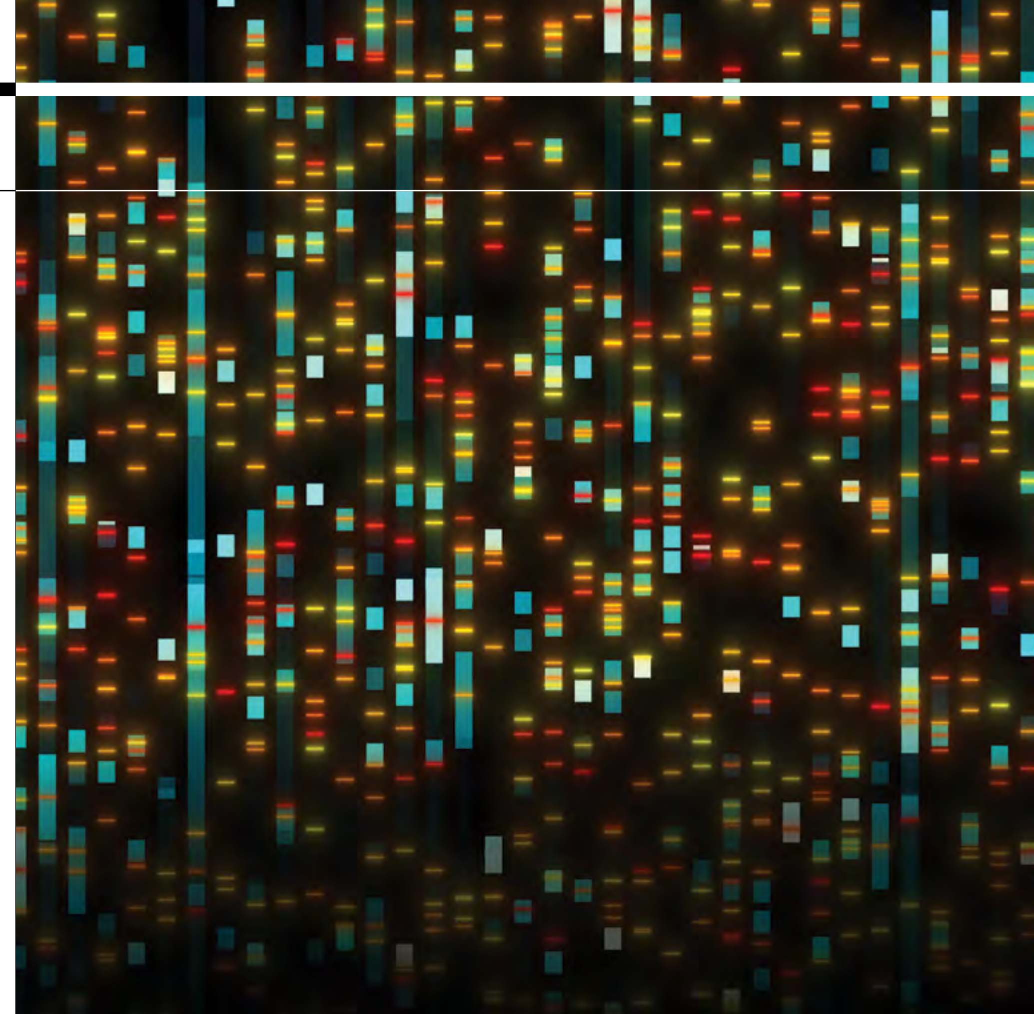

QUANTUM TECHNOLOGY TAKES ON CANCER

No longer exclusive to the halls of universities and big corporations, quantum technology is inching its way closer to public visibility through the privately-owned Cleveland Clinic in Ohio, United States. The hospital's Global Center for Pathogen Research and Human Health will be the first external facility to deploy Big Blue's Quantum System One, as part of a ten-year partnership with IT giant International Business Machines Corporation (IBM).

As IBM's highest performing quantum technology to date, the system has up to 18 entangled qubits and averages a two-qubit error rate of less than two percent, leading to game-changing speed and accuracy in processing data. With this latest installation, quantum computing is poised to accelerate the clinic's healthcare research, building on top of previous collaborations including a machine learning platform for cancer research.

This time, anticipated use cases for the new quantum solutions include analyzing genomic data, identifying potential targets for drugs and developing new health technologies for diagnostics and treatment. Given the fast-evolving nature of quantum technology, the two parties are already planning to deploy IBM's next quantum system at Cleveland Clinic in 2023.

SAUDI ARAMCO'S DAMMAM BREAKS INTO TOP TEN

Oil and gas corporation Saudi Aramco, in partnership with IT firm Saudi Telecom Company (STC), has deployed the 10^{th} most powerful supercomputer in the world: Dammam 7. In a list that seldom sees changes at the top, Dammam 7's historic entry into the rankings comes off the back of its 22.4 LINPACK petaFLOPS of computing power delivered by 6,248 Intel Xeon Gold CPUs.

The supercomputer, which can hold more than 506 terabytes of data, was manufactured by HPE Cray, Solutions by STC and Dhahran Techno Valley. It is the latest addition to Saudi Aramco's growing arsenal of supercomputing resources, four of which also rank among the TOP500 fastest computers. Through imaging and deep learning techniques, these systems are enhancing the efficiency of oil and gas discovery and recovery.

"New discoveries and enhanced recoveries are crucial to both ensuring the availability of adequate supply to meet the demand for energy and to cut costs while boosting productivity," said Amin Nasser, Saudi Aramco president and CEO. "Continuous investment in innovative technology is an essential enabler of our company's long-term growth strategy."

Super Snapshot

SCALING UP SINGAPORE'S SUPERCOMPUTING STRENGTH

ASPIRE1
Launch date
2016
Top speed
1 petaFLOPS
Compute cores
30,912 x Intel Xeon CPUs
Accelerators
128 NVIDIA Tesla GPUs

HPE Cray EX
Launch date
Early 2022
Top speed
10 petaFLOPS
Compute cores
>100,000 AMD EPYC CPUs
Accelerators
352 NVIDIA A100 Tensor Core GPUs

THE NEW MACHINE AT THE NATIONAL SUPERCOMPUTING CENTRE SINGAPORE (NSCC), SLATED TO COME ONLINE IN 2022, WILL BE EIGHT TIMES MORE POWERFUL THAN ITS PREDECESSOR.

Since its launch in 2016, Singapore's ASPIRE1 supercomputer has been in hot demand, reliably operating at over 90 percent capacity. As part of a S$200 million investment into supercomputing infrastructure, the country has set aside S$40 million for a new HPE Cray EX system designed to support exascale-class computing.

New features

- Expanded storage for complex workloads
- Purpose-built networking for congestion control
- Software suite optimized for both HPC and AI applications

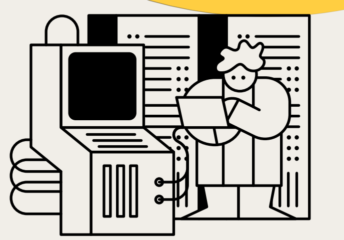

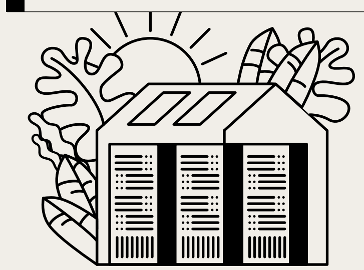

Meaner and greener

- Housed in a data center designed for sustainability
- Warm water-cooled, the first of its kind in the tropics
- Increased energy efficiency and power density

In the service of sustainability

- Weather forecasting and climate studies for Singapore and Southeast Asia
- Environmental modeling for urban sustainability
- Designing materials for the low-carbon and circular economy

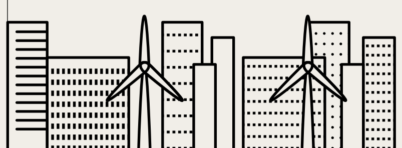

POWERING RESEARCH TO COOL SINGAPORE

with the development of the Digital Urban Climate Twin by the Cooling Singapore team at the Singapore-ETH Centre

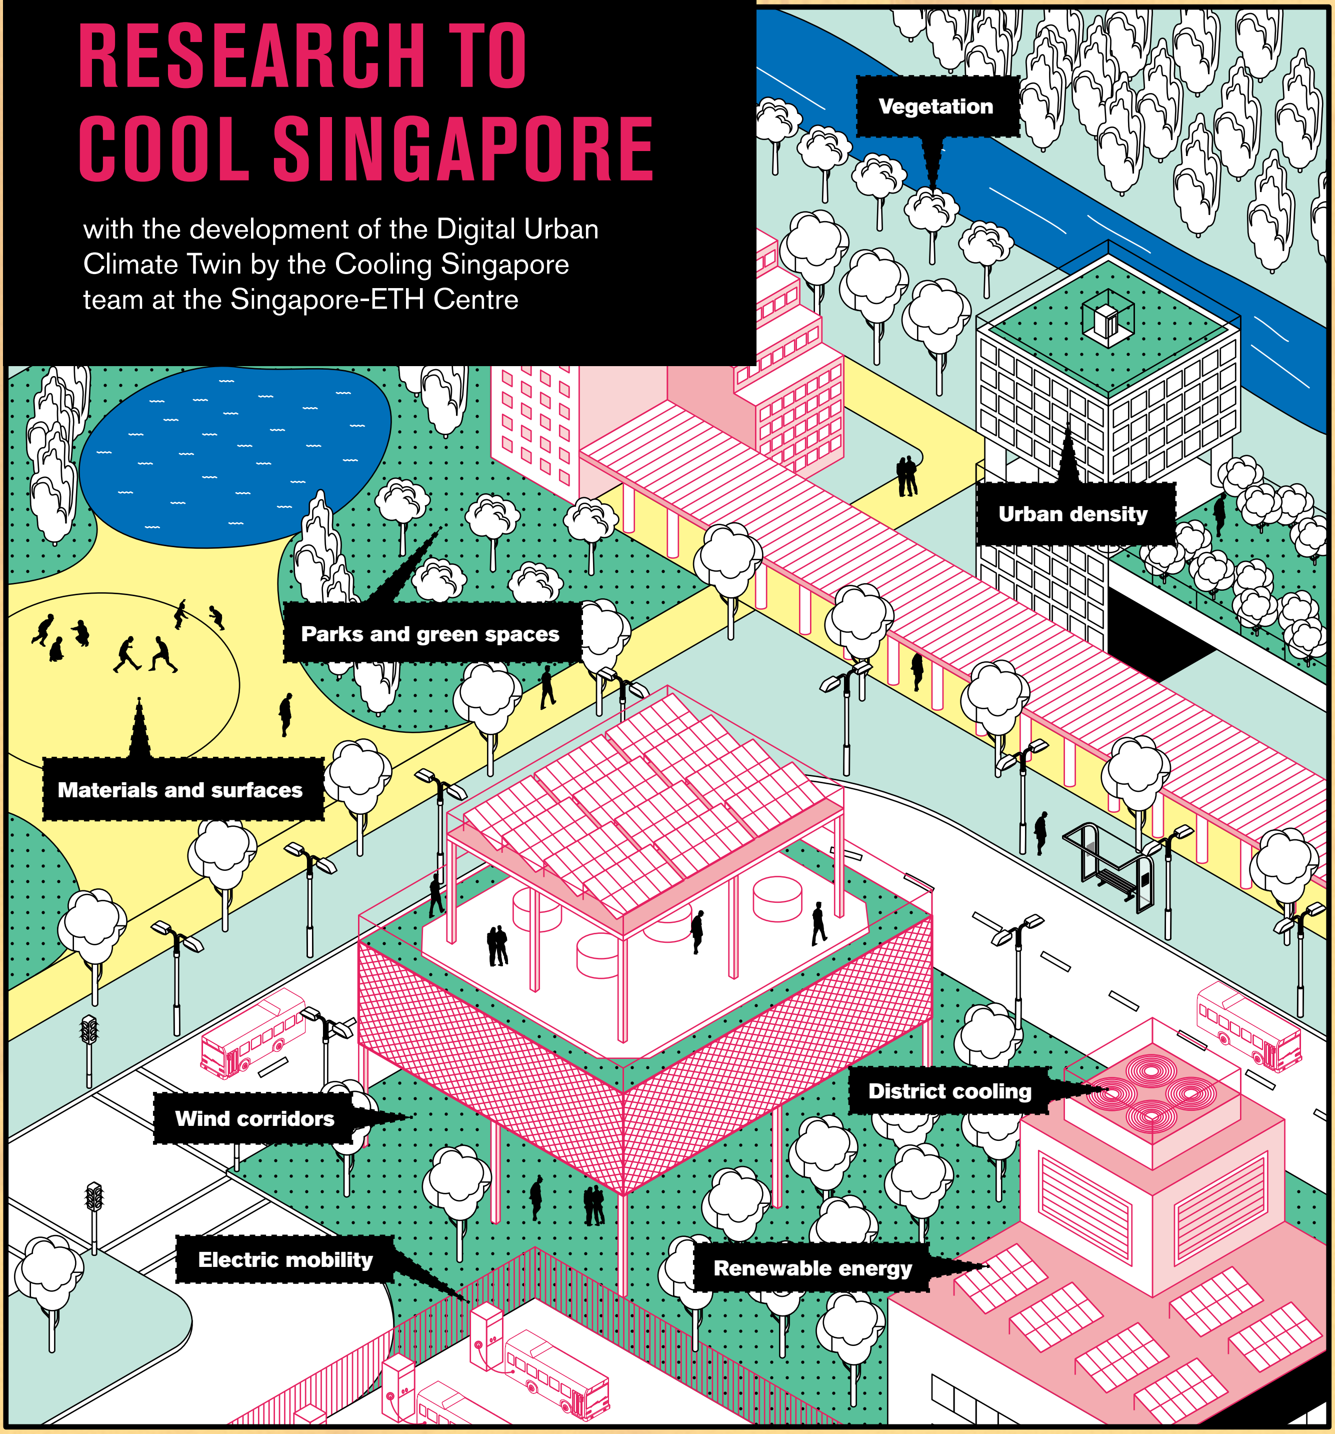

A collaboration powered by

(SEC) SINGAPORE-ETH CENTRE

SG NSCC SINGAPORE National Supercomputing Centre

www.coolingsingapore.sg