


ASIAN SCIENTIST

Issue 08
July 2020

SUPERCOMPUTING ASIA

MAPPING THE BRAIN'S MYSTERIES

HARNESSING HPC
TO UNDERSTAND
THE BRAIN



JAPAN'S FUGAKU

ON TOP OF
THE WORLD

REVVING UP
COVID-19 RESEARCH

SUPERCHARGING
HEALTHCARE
INNOVATION

EL CAPITAN

On track to be the world's fastest supercomputer !

WHAT COULD YOU DO WITH THE WORLD'S FASTEST SUPERCOMPUTER?

If you're Lawrence Livermore National Laboratory (LLNL), you'll make the world safer.

LLNL focuses on



National security



Addressing issues from **nuclear stockpile stewardship** to **threats posed to the environment**



Public health

SPEED

Their new **HPE CRAY SUPERCOMPUTER** will help them do it at an anticipated

2 QUINTILLION CALCULATIONS PER SECOND.

Just how fast is **2 QUINTILLION?**

IT'S **FAST — 18 ZEROS FAST!**

2 quintillion = 2 exaflops = 2,000 petaflops
(2,000,000,000,000,000,000)



IF ALL **7.7 BILLION**

people on Earth each completed one calculation per second, it would take over **8 YEARS** to do what LLNL's **EL CAPITAN** will be capable of doing in **1 SECOND.**

POWER

EL CAPITAN will be **MORE POWERFUL** than the current

TOP 200 fastest supercomputers in the world **COMBINED**¹ and **10X**

faster than the world's current fastest system.²



IMPACT

Exascale supercomputing delivers advanced capabilities in **MODELING, SIMULATION,** and **ARTIFICIAL INTELLIGENCE** that will impact every human life for decades to come.

Systems like **EL CAPITAN** will help change outcomes for critical issues like **CLIMATE CHANGE** which already costs **150,000 LIVES** and billions of dollars annually³ and **CANCER** which is expected to increase to **23.6 MILLION** new diagnosed cases per year by 2030.⁴

LEARN MORE AT [CRAY.COM/ELCAPITAN](https://www.cray.com/elcapitan)

^{1,2} Based on Rpeak numbers from Top500 system list, November 2019

³ [scientificamerican.com/article/global-warming-and-health/](https://www.scientificamerican.com/article/global-warming-and-health/)

⁴ [cancer.gov/about-cancer/understanding/statistics](https://www.cancer.gov/about-cancer/understanding/statistics)

© Copyright 2020 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein. All third-party marks are property of their respective owners.

CONTENTS

Issue 08
July 2020

COVER STORY

p. 24

Mapping The Brain's Mysteries

Harnessing HPC to understand the brain

FEATURES

p. 32

Getting A Headstart On HPC

HPC at institutes of higher learning

p. 6

Digital Dispatch

Supercomputing news from around the world

p. 10

Supercharging Healthcare Innovation

5 ways supercomputers are shaping the future of healthcare

p. 36


Business Bytes

The latest industry moves

p. 38

Super Snapshot

Japan's Fugaku—on top of the world



FEATURES

p. 16

REVVING UP COVID-19 RESEARCH

Supercomputing takes on SARS-CoV-2

EDITOR'S NOTE

When we first dreamt up this issue of *Supercomputing Asia* more than six months ago, we had no idea how important our chosen focus on healthcare and biomedical applications of high performance computing would become. We planned to shine a spotlight on the researchers across Asia who are using supercomputers to take on a monumental task: creating a high-resolution map of the trillions of neurons that make up the human brain (*Mapping the Brain's Mysteries*, p. 24).

As soon as the novel coronavirus SARS-CoV-2 began spreading across the world, however, it became clear that supercomputing would be essential for finding a way forward. Supercomputers not only enable us to understand the virus and accelerate the discovery of new drugs, but they also help to inform decisions about non-medical interventions like masks and indoor ventilation (*Revvng Up COVID-19 Research*, p. 16).

Driven by the urgency of research addressing the pandemic, Japan has brought part of their post-K supercomputer, Fugaku, online ahead of schedule. Already put to use on COVID-19-related workloads, the impressive machine has debuted as the number one system on the TOP 500 list as well as the HPCG, HPL-AI and Graph500 rankings, making it the first machine to hold top positions across all four fields of measurement. Find out more on p. 38!



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

- /asianscientist
- @asianscientist
- Asian Scientist Magazine
- 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

EDITORIAL EXECUTIVE

Kamila Navarro

CONTRIBUTORS

Eugene Low
Li Lidao
Tim Hornyak

ART DIRECTOR

Lam Oi Keat

SENIOR DESIGNER

Shelly Liew

SALES & MARKETING

Samantha Yeap
Sammi Ting
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



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

Brands under Wildtype Media Group include the flagship *Asian Scientist Magazine* and *Supercomputing Asia*, award-winning titles available in print and online.

www.wildtype.media

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

Supercomputing Asia, MCI (P) 025/04/2020, is published by Wildtype Media Group Pte Ltd. Printed by KHL Printing Co Pte Ltd. All rights reserved. No part of this publication is to be reproduced, stored, transmitted, digitally or otherwise, without the prior consent of the publisher. The information contained herein is accurate at time of printing. Changes may have occurred since this magazine went to print. Supercomputing Asia content is provided 'as is' without warranty of any kind. Supercomputing Asia excludes all warranties, either expressed or implied. In no event shall Wildtype Media Group Pte Ltd and its editors be held liable for any damages, loss, injury, or inconvenience, arising, directly or indirectly, in connection with the contents of the magazine. Any persons relying on Supercomputing Asia content shall do so at their own risk. Advertising material submitted by third parties in Supercomputing Asia is the sole responsibility of each individual advertiser. We accept no responsibility for the content of advertising material, including, without limitation, any error, omission or inaccuracy therein.

ASIAN SCIENTIST

MEET

ASIA'S


A video series that profiles innovators


CHANGE

and leaders in science and technology.

MAKERS

Subscribe to Asian Scientist on YouTube to watch the full series!






DISTINGUISHING COVID-19 PATIENTS FROM THE START

Chinese clinicians are using artificial intelligence (AI) and high performance computing to quickly distinguish between patients with COVID-19 or common viral pneumonia. The AI-powered system will assist clinicians in the early screening of suspected coronavirus cases, as images of viral pneumonia tend to look similar.

Jointly developed by the National Supercomputer Center, Tianjin Medical University Cancer Hospital and Institute, and other domestic hospitals, the system analyzes chest scans generated from computed tomography (CT). This was done with a convolutional neural network, a deep learning algorithm that is typically applied in image classification. By training the neural network on previously obtained COVID-19 CT scans, the algorithm managed to extract hallmarks of coronavirus infection for use as a diagnostic reference.

Now running on China's first petaFLOP supercomputer, Tianhe-1, the AI system can easily breeze through 300 images and give a diagnosis in about ten seconds with a classification accuracy of over 80 percent. For comparison, it would take an experienced physician about 15 minutes to go through the same number of CT scans. A shorter diagnosis time can lead to quicker action—something crucial given that COVID-19 patients are said to be most infectious during the early stages.

STREAMLINING THE SEARCH FOR CORONAVIRUS ANTIBODIES



A computational estimate of the structure of S230, an antibody known to have neutralizing activity against the severe acute respiratory syndrome (SARS) coronavirus

Credit: Lawrence Livermore National Laboratory

At the Lawrence Livermore National Laboratory (LLNL) in the United States, scientists are harnessing a powerful combination of AI, bioinformatics and supercomputing to accelerate COVID-19 drug discovery on all fronts.

Using a suite of structural modeling and analysis tools, LLNL senior computer scientist Dr. Adam Zemla was able to predict the coronavirus' 3D structure. Next, LLNL data scientists Dr. Daniel Faissol and Dr. Thomas Desautels leveraged a pair of HPC clusters to simulate the molecular interactions between known SARS antibodies and the virus. Through their platform, the team identified sites on the SARS antibodies that could be modified so that they could bind instead to SARS-CoV-2.

Now, another group at the LLNL is utilizing the Intel Xeon-based Quartz supercomputing cluster, which has a peak performance of 3.25 petaFLOPS, to screen over 26 million molecules against four of the virus' key protein sites. "Using the computational tools and data that we [previously] created... we were able to computationally screen these molecules quickly and at a large scale," said project leader Dr. Felice Lightstone, a senior scientist at LLNL. "This is the first step towards finding a new antiviral."

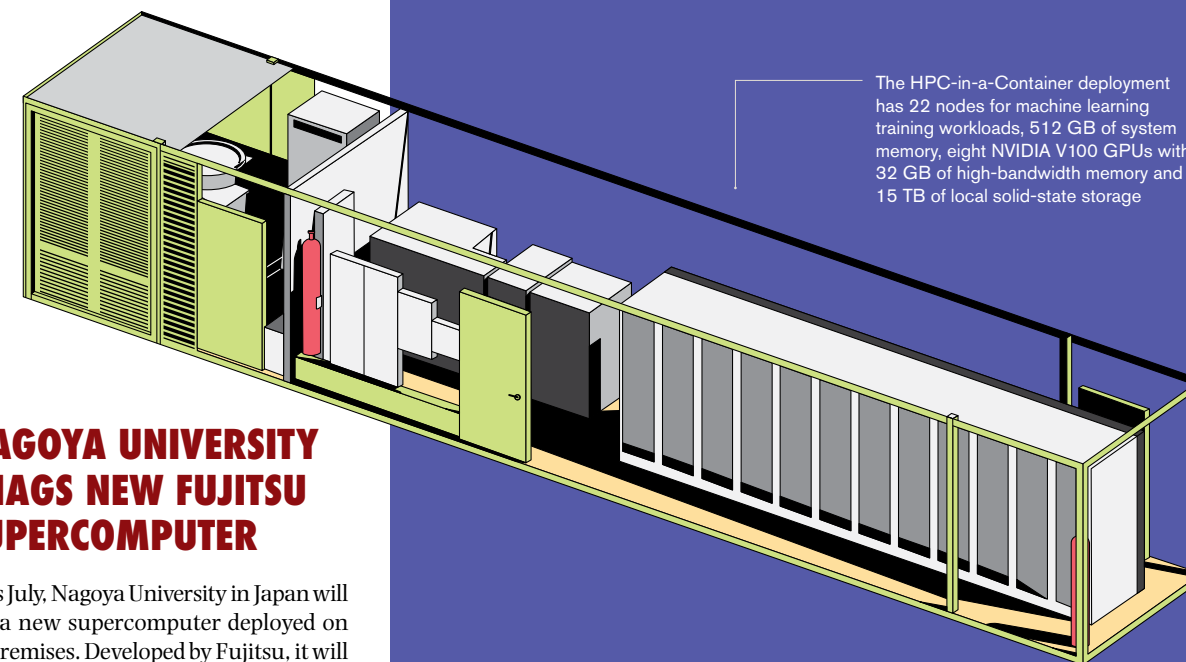
NAGOYA UNIVERSITY SNAGS NEW FUJITSU SUPERCOMPUTER

This July, Nagoya University in Japan will see a new supercomputer deployed on its premises. Developed by Fujitsu, it will be the first commercial supercomputer powered by the ARM-based A64FX technology developed for RIKEN's Fugaku supercomputer (see infographic on p. 38).

The new system will consist of 2,304 Fujitsu PRIMEHPC FX1000 nodes, delivering a theoretical peak performance of 7.782 petaFLOPS and a total main memory capacity of 72 terabytes. The incoming supercomputer represents a significant upgrade from its current system, a Fujitsu PRIMEHPC FX100 cluster of 2,880 nodes with a peak performance of 3.2 petaFLOPS.

This will provide a much-needed boost in computing capacity, given the increasing numbers of researchers engaged in data science, modeling and simulation nationwide. With the new system, Nagoya University will be able to explore applications ranging from drug design, numerical simulations of extreme weather events like supertyphoons and tsunamis, and even AI in autonomous driving.

The new supercomputer will be made available not just to Nagoya University staff and students, but also to researchers in universities and research institutes across Japan.



The HPC-in-a-Container deployment has 22 nodes for machine learning training workloads, 512 GB of system memory, eight NVIDIA V100 GPUs with 32 GB of high-bandwidth memory and 15 TB of local solid-state storage

PRESENTING THE SHIPPING-CONTAINER SUPERCOMPUTER

The United States Department of Defense (DoD) has deployed a US\$12 million six-petaFLOPS supercomputer in a surprising place: a mobile shipping container. Built as part of the DoD's High Performance Computing Modernization Program, the unusual supercomputer includes 1.3 petabytes of solid-state storage and will be used for both training and inference workloads.

The shipping-container supercomputer consists of 22 compute nodes for machine learning training and 128 for inference, each with GPU acceleration and two IBM Power9 processors. For training, it relies on 22 IBM AC922 servers

each equipped with six NVIDIA V100 GPUs. It also features multiple IBM IC922 servers focused on AI inference, with an additional six NVIDIA T4 GPUs on each server for AI inference acceleration.

According to the DoD, the supercomputer is designed to be deployed at the tactical AI edge—meaning that data is collected and processed at the same location. This will enable its application for military purposes that would have been impossible with fixed supercomputers. Currently, opportunities for deployment to remote locations are being explored.


Digital Dispatch

MAPPING THE HUMAN MIND

Scientists from Singapore, Japan, South Korea and Taiwan are collaborating to produce a first-of-its-kind, ultra-high-resolution 3D comprehensive map of the human brain's neural network.

Using synchrotrons—extremely powerful X-ray sources—scientists aim to trace the brain's intricate networks. The project will link the synchrotron facilities in the Asia Pacific region under a collaboration called Synchrotron for Neuroscience—Asia Pacific Strategic Enterprise (SYNAPSE). For more details, see our cover story on p. 24!

Each participating facility will work on and image a portion of the same brain. As mapping a human brain will generate a huge amount of data, a high performance computing network will be developed for the project. Singapore's petascale supercomputing resources at the



Credit: SYNAPSE

National Supercomputing Centre (NSCC) Singapore will serve as the data hub of SYNAPSE, linking all of the partners via its high-speed international network connections.

"The images captured by SYNAPSE will form an extensive human brain map," said Low Chian Ming, an associate professor at the National University of Singapore who co-founded the SYNAPSE consortium. "The map will show how neurons are connected and how they interact to result in cognition and intelligence. Our findings could potentially contribute to effective treatments for increasingly important neurodegenerative pathologies such as Alzheimer's disease and other forms of dementia."

ISC20 CLUSTER COMPETITION GOES VIRTUAL

From conferences to student competitions, practically all in-person events have been put on hold due to COVID-19. Still, the show must go on, and so the ISC Student Cluster Competition took place virtually for the first time in its history.

Co-organized by the HPC-AI Advisory Council and ISC, the competition is an annual event where undergraduate student teams assemble HPC clusters and race to complete scientific benchmarks and applications the fastest. Fourteen teams from 11 different countries participated this year.

As a result, the competition was markedly different from its previous iterations. For instance, the teams had to remotely drive the exact same cluster provided by the National Supercomputing Centre

(NSCC) Singapore—leveling the playing field in terms of hardware. It was the largest system many of them have ever used, the cluster is composed of two nodes: the first being a typical Xeon-based system, while the second node is an NVIDIA DGX-1 box with eight V100 GPUs.

Because all teams used the same cluster, the organizers prepared real-world HPC test applications instead of the usual benchmarks. These applications include those being used by real life scientists to accelerate the search for COVID-19 treatments. The event took place from June 1–24, 2020, with a team from the University of Science and Technology of China 'walking' away with the first prize at an live online awards ceremony held on the last day.

PEERING INTO THE HORSESHOE CRAB'S EVOLUTIONARY HISTORY

Horseshoe crabs have existed for nearly 500 million years, but they've been threatened by human over-exploitation as well as habitat loss.

"An important prerequisite for formulating effective conservation strategies is an understanding of their genetic diversity and population structure," explained Dr. Byrappa Venkatesh of the Institute of Molecular and Cell Biology at the Agency for Science, Technology and Research, Singapore. "However, at present, there is no information about the population genetics of horseshoe crabs in Singapore."

Venkatesh and his team generated 225 GB worth of genomic reads, along with 200 GB of shorter reads for error-correction. The team then leveraged the National Supercomputing Centre (NSCC) Singapore's resources to generate a chromosome-level assembly of the horseshoe crab genome from scratch using 50 CPU nodes and large-memory nodes with up to 7 TB of memory. They also performed gene annotation and comparative analyses.

"Our study has shown a marked decline in the horseshoe crab population size approximately 60,000 years ago coinciding with the onset of the recent ice age. The species has not recovered from this decline," remarked Venkatesh. Given this, he suggests that horseshoe crab stocks could be managed by prioritizing habitat conservation, prohibiting their harvest and protecting their breeding grounds.



WHAT'S UP!

SAVE THE DATE: SUPERCOMPUTINGASIA 2021

SupercomputingAsia (SCA), Asia's premier gathering of the region's best minds in HPC, will be returning to Singapore in 2021 after a brief hiatus due to the coronavirus pandemic. In 2019, the conference was attended by more than 700 delegates from over 20 countries across Asia, America and Europe.

Organized by the National Supercomputing Centre (NSCC) Singapore, SCA21 will take place from March 1–4, 2021, at the Suntec Singapore Convention and Exhibition Centre. As always, you can look forward to an impressive lineup of keynote addresses from the world's top HPC experts and thought leaders as well as a jam-packed calendar of research presentations, workshops and other exciting HPC activities. Save the date—more details will be coming your way soon!

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

WHAT

SupercomputingAsia 2021

WHEN

MARCH 1–4, 2021

WHERE

Singapore

STILL ON: SUPERCOMPUTING 2020

It's still all systems go for Supercomputing 2020 (SC20), which will be held in Atlanta, Georgia from November 15–20, 2020. Also known as the International Conference for High Performance Computing, Networking, Storage and Analysis, SC20 is the world's leading conference on all things HPC.

This year marks the launch of a new program known as HPC Immersion. Designed to introduce students to HPC careers, the program will target traditionally underrepresented students majoring in electrical engineering, computer engineering and computer engineering.

SC20 will also feature three full days of dedicated State of the Practice Talks, meant as an avenue for the discussion of practical improvements in supercomputer operation and infrastructure. These presentations will broadly cover topics like system management, networking and data management. Another thing to look forward to is HPC in the City, an initiative to engage the HPC community in Atlanta to address relevant local problems that could be solved through HPC.

Currently, the conference is planned to be a hybrid event, with both on-site events and an enhanced virtual conference. Further updates will be given closer to the date.

For more information, visit <https://www.sc20.supercomputing.org>

WHAT

SC20

WHEN

November 15–20, 2020

WHERE

Atlanta, Georgia, US

Supercharging Healthcare Innovation

From bench to bedside, high performance computers are running simulations, crunching big data and suggesting medical treatments to make healthcare more effective and affordable.


By **Li Lidao**

As the COVID-19 pandemic rages across the globe, medical professionals, scientists and public health officials are in a race against time to curb the spread of the novel coronavirus.

While we may need them now, the development of new diagnostic tools, vaccines and therapeutics remains complex and may take longer than we can wait for. This is where supercomputers, well known for their ability to rapidly process large amounts of data, can help to accelerate the process. After all, the top 500 supercomputers in the world can perform more than 1 quadrillion—that's 1 with 15 trailing zeros—operations per second on average.

Beyond tackling the immediate challenges of infectious diseases, supercomputers can also be applied to other aspects of healthcare. From laboratory based medical research to clinical practice, here's how supercomputers are augmenting human abilities and helping medical professionals deliver personalized care.

5 ways supercomputers
are shaping the future
of healthcare



Discovering new therapeutics


Supercomputers that simulate how complex biological molecules behave have emerged as a useful tool in responding to the COVID-19 pandemic.

Armed with a peak performance of 1.3 petaFLOPS, the MDGRAPE-4A supercomputer at Japan's RIKEN completed a simulation of the protease protein involved in the replication of the SARS-CoV-2 virus on March 17, 2020.

More than being just a static image, the simulation showed how the 2,416 atoms making up the protease protein move and wobble around in solution, allowing scientists to screen for potential antiviral compounds that can block it.

Such a trial-and-error process could be performed virtually as well, with supercomputers iterating through a much wider range of compounds than what is physically possible. Although not specifically applied to COVID-19 research, a software framework developed at the National Supercomputer Center in Guangzhou, China, trawled through ten million molecules in a trial run, taking just 22.31 hours using the Tianhe-2 supercomputer, which boasts a peak performance of close to 34 petaFLOPS.

Simulating virtual organs



In the past, we studied cadavers to unravel the inner workings of the human body. While gaining access to human organs is often challenging, computer simulations allow researchers to safely experiment with virtual organs that are programmed to respond in the same way.

To reproduce an organ's complex physiology, researchers need to take into account multiple factors ranging from biochemical reactions at the subcellular level, all the way up to mechanical behavior at the organ level. Such extensive simulation can be accomplished with the help of a supercomputer.

In Japan, the T2K Open Supercomputer at the University of Tokyo was used to simulate the heart and reproduce the expected side effects of a heart problem known as arrhythmia when certain drugs were administered. Dubbed UT-Heart, the simulation can be used to screen for new medications for adverse side effects before more costly laboratory and clinical testing are conducted.

In addition, UT-Heart can even be customized based on each patient's medical data. The team used it to predict the effectiveness of a pacemaker on each individual, paving the way for doctors to develop personalized treatment plans.

Rapid genome sequencing

In medicine, there is rarely a one-size-fit-all approach—two individuals with the same disease may show different symptoms and respond very differently to the same treatment. For example, the SARS-CoV-2 virus leads to severe pneumonia in some cases, but only a mild cough in others.

Researchers think that this is in part due to subtle variations in individual genomes, and mapping them out could suggest why certain patients are less susceptible to infections, possibly pointing the way to more effective treatments and vaccines.

Chinese company BGI Genomics, in partnership with Intel and Lenovo, has thus sought to sequence the genome of a large number of COVID-19 patients to find out why. However, genome sequencing is a time-consuming task—the first human genome took over a decade to decode and it now typically takes about a week, even with modern technology.

To give high performance computers an extra boost, Lenovo fine-tuned both hardware and software, designing an architecture known as the Genomics Optimization and Scalability Tool (GOAST) that is able to sequence an entire genome in just 5.5 hours.

Once a vaccine has been developed, gene sequencing tools could also allow scientists to predict which subpopulation of patients it would work best in.



Reading medical images

2019 was the first time there were more elderly above 65 than young people below five. As aging populations place increasing demands on healthcare, high performance computers could be called in to help.




One possibility lies in the use of artificial intelligence (AI) to help doctors analyze medical images and make diagnoses. As a job that requires many years of training in humans, it is no easier for machines. Medical data such as those from computerized tomography (CT) scans are three-dimensional, requiring huge amounts of computational power to train an AI software to recognize them.

The process, however, can be sped up by using supercomputers. A team from The University of Tokyo in Japan used the Reedbush-H supercomputer, which has a peak performance of 800 teraFLOPS, to train an AI

that could automatically detect nodules (which might be cancerous) from lung CT images.

While certain training steps are typically carried out sequentially, the team redesigned the protocol to carry them out on multiple parallel GPUs simultaneously, drastically reducing the time required to train the AI from around 105 hours to 40 hours.

Once trained, the AI could assist doctors in making a diagnosis by, for example, identifying small nodules that might have been missed.



Suggesting medical treatments

While we are still far from robot doctors taking over hospitals, there are early signs that medical AI could assist doctors by providing treatment recommendations. It may even do a better job than doctors as the AI is able to scan through more medical literature than what a human could ever read.

One key challenge lies in teaching computers how to interpret information in words, in addition to numbers. IBM's Watson for Oncology supercomputer is one example to have emerged.

When trialed at the Manipal Comprehensive Cancer Center in India, Watson read through patient records and proposed treatment courses that were in line with a physician's recommendation in 73 percent of 638 breast cancer cases. The results of similar trials were 83 percent in Thailand and 49 percent in South Korea.

Clearly, IBM's supercomputer is still a work in progress, but applications for high performance computing in healthcare are likely to become more common in the future. To meet increasing demand, Taiwanese company Infortrend Technology Inc. developed a data storage system named EonStor CS to facilitate the high-speed sharing of medical files. To keep it 'future proof,' it was designed to be easily scaled up to accommodate a growing amount of data in a cost-effective manner.

REWINDING UP COVID-19 RESEARCH

**Supercomputing
takes on SARS-CoV-2**

With the world racing to fight the coronavirus, Asia's most powerful supercomputers are entering the fray.


By **Kamila Navarro**

In the classic military treatise *The Art of War*, Sun Tzu emphasizes that a deep understanding of the enemy holds the key to their defeat. Though COVID-19's battlegrounds look vastly different from ancient China—think sterile laboratories and hospital rooms instead of rugged terrain—the same principle still applies in the global fight against the coronavirus.

As soon as the outbreak's extent became apparent, researchers worldwide were quick to react. Just a month after the first reported case in Wuhan, Chinese scientists had already released the genome of SARS-CoV-2—the virus behind COVID-19—allowing countries worldwide to swiftly create diagnostic kits.

Successfully developing interventions for COVID-19, however, involves a two-pronged approach: first, analyzing the biological molecules that equip the coronavirus with its deadly abilities. The molecules could then be used as potential targets for drugs or even vaccines. Second, making sense of the virus' knack to stealthily spread in various environments. The thing is, these approaches take time. For instance, drug discovery and vaccine development both take years, if not decades.

With no time left to lose, scientists are hoping to accelerate the process with high-powered help. Across Asia, countries like Singapore, Japan and South Korea have launched special calls for proposals that leverage some of the world's fastest supercomputers to fight the pandemic. Here's a look at a few ongoing projects that could turn the tide in the battle against COVID-19.



Structural model of the SARS-CoV-2 spike protein

“WE ARE SEEKING ARTIFICIAL INTELLIGENCE-ASSISTED APPROACHES TO PREDICT THE MAIN PROTEASE'S STRUCTURE, TAKING INTO CONSIDERATION ANY MUTATIONS.”

DR. CHENG YUAN

SENIOR SCIENTIST, A*STAR INSTITUTE OF HIGH PERFORMANCE COMPUTING

FINDING THE WEAK SPOTS

Over at A*STAR's Institute of High Performance Computing in Singapore, Dr. Cheng Yuan is leading a team to investigate the structure of SARS-CoV-2's main protease. Responsible for cutting precursors of viral proteins into functional pieces, the main protease plays a critical role in mediating viral replication and transcription, said Cheng. However, predicting a protein's structure based on its sequence has always been a significant challenge. In addition, it's entirely possible for the main protease to mutate—further complicating structural predictions.

“We are seeking artificial intelligence-assisted approaches to predict the main protease's structure, taking into consideration any mutations,” explained Cheng. Their approach will combine machine learning with multiscale modeling—a simulation strategy that simultaneously considers models at different scales of complexity. Given that proteins have four structural levels, she anticipates their experiments to be computationally intensive.

“The process is very demanding because of the large size of the data set and molecular simulations,” noted Cheng. Accordingly, her team will be tapping upon the powerful computational resources of Singapore's National Supercomputing Centre (NSCC). Specifically, they'll be leveraging the flagship ASPIRE1 petascale supercomputer as well as an eight-GPU NVIDIA DGX-1 AI system with V100 cards and 13 petabytes of high performance storage.

Delineating the main protease's structure would provide a deeper understanding of how it functions, and more intriguingly, give insights into its active site. When the active site is bound by an inhibitor, the protease is unable to function, interrupting viral replication and transcription. Just like striking an enemy's weak spot in battle, knowing the main protease's structure should aid in the design of drugs that inhibit its function.

Another SARS-CoV-2 protein being closely studied by researchers is the spike protein, recognizable as the distinctive spikes that dot the surface of the virus. The coronavirus uses the spike protein to bind to and invade human cells, through a receptor called ACE2. Similar to his counterparts in Singapore, Dr. Yuji Sugita from Japan's RIKEN is seeking to predict the spike protein's structure by simulating the way its atoms and molecules dynamically move over time. This technique, known as molecular dynamics (MD) simulation, would allow his team to discover structures that cannot be obtained through conventional means.

To achieve this, they'll be running MD simulations on RIKEN's Fugaku supercomputer, which is still in the process of installation. Despite this, even the partially-available Fugaku is expected to run the simulations at a speed 125 times faster than its predecessor, the K supercomputer. By pinning down the complex structure of the spike protein, Sugita's findings could help inform the development of drugs that block the interaction between the spike protein and the ACE2 receptor—preventing the virus from binding to the cells in the first place.

DRIVING DRUG DISCOVERY

Due to the pressing need for a COVID-19 treatment, alternative tactics are being used to shorten the time frame for drug discovery. To save time, rather than finding drugs completely new to science, scientists are repurposing drugs that have already been approved. Indeed, all treatments currently being tested in the World Health Organization's global Solidarity trial are approved for use in other indications such as HIV or malaria.

To speed up the discovery process even further, a research team led by Professor Seo Sangjae of the Korea Institute of Science and Technology Information (KISTI) tapped on the world's 14th fastest supercomputer, Nurion, to computationally screen thousands of drugs that could be given a second chance. Starting with a pool of almost 20,000 compounds sourced from the SWEETLEAD library and ChEMBL database, Seo and his team systematically evaluated the binding affinity of these compounds to the main protease of SARS-CoV-2 through a technique called molecular docking.

Comparable to finding the right key to a lock, the team then calculated the docking score by assessing which orientations of the compounds best fit the main protease's active site and measuring the binding strength between the two molecules. Among those with the highest docking scores, the team chose 43 compounds to be investigated further using MD

simulations. This motley set of compounds included antiviral drugs, antibiotics for pneumonia, vitamins, and drugs in clinical trials such as remdesivir and hydroxychloroquine.

"To identify drugs, it's important to understand their interactions with enzymes," remarked Seo. "Unlike other research projects that only performed molecular docking, we improved the accuracy of our results by conducting MD simulations as well." Publishing their preliminary results on *ChemRxiv* in early April, his team identified eight promising COVID-19 drug candidates that were all antivirals for either hepatitis C or HIV.


Incredibly, the team crunched all these calculations and simulations on Nurion in just one week. "Each MD simulation was computationally demanding. To save time, we even had to run 43 drugs simultaneously. Harnessing supercomputing resources was the only way we could do this," explained Seo. Considering that Nurion is equipped with Intel Xeon Phi 7250 and 8,305 nodes, a study of a similar scale would require at least 200 days to finish on a personal computer.

Moving forward, Seo shared that the 43 drugs previously chosen for MD simulations are currently being experimentally tested in collaboration with KISTI's partners. The experimental results will be compared with the MD simulations to help improve the computational process. At the same time, the team is now looking to make their docking calculations even more accurate by adding artificial intelligence into the mix.

“UNLIKE OTHER RESEARCH PROJECTS THAT ONLY PERFORMED MOLECULAR DOCKING, WE IMPROVED THE ACCURACY OF OUR RESULTS BY CONDUCTING MOLECULAR DYNAMICS SIMULATIONS AS WELL.”

PROFESSOR SEO SANGJAE

KOREA INSTITUTE OF SCIENCE AND TECHNOLOGY INFORMATION (KISTI)



Credit: Makoto Tsubokura/RIKEN

PROFILING A KILLER

Aside from poring over SARS-CoV-2's structure, it's also equally important to consider how the virus is transmitted from person to person. After all, though the pandemic is nowhere close to ending, countries are already cautiously emerging from their respective lockdowns. As more and more people begin their much-awaited return to the outside world, it is inevitable that we'll be seeing fairly crowded spaces once again.

Unfortunately, crowds are also prime breeding grounds for the coronavirus. Human transmission is thought to mainly occur via droplet spread, when an infected individual releases thousands of droplets into the surrounding air via coughing, sneezing or even talking. But in a sinister twist, it's been suggested that these droplets could be aerosolized into even smaller particles that could spread over larger distances and linger in the air.

Aerosol transmission poses a problem, especially in enclosed environments like public transport, most workplaces, classrooms and restaurants. To facilitate the safe resumption of activities that take place in these settings, Professor Makoto Tsubokura of RIKEN and Kobe University is looking to simulate the scattering of virus droplets under varying conditions on the Fugaku supercomputer (see our infographic on p. 38).

Their project will consist of three steps, explained Tsubokura: first, assessing the infection risk of droplet and aerosol transmission in daily scenarios, including commuting and offices. Based on their results, his team then hopes to propose immediate countermeasures that could reduce infection risk—ranging from opening or closing windows to the strategic placement of partitions. Finally, they seek to suggest long-term measures in the fight against COVID-19, such as improved air conditioning or ventilation systems.

“We are required to produce many results in a very short period of time,” said Tsubokura. “In addition, in cases like train simulations, we have to consider very crowded cabins with more than 200 passengers, running at 80 kilometers per hour.” Given that the team has to evaluate the ventilation effect of open windows in such complicated scenarios, Fugaku’s massive computational resources are required for their simulations. His team’s academic and industry partners will then go on to experimentally validate their results.

From the simulations, Tsubokura is also aiming to create accessible animations of the virus’ droplet and aerosol spread so that the public can easily understand the risk of infection and the need for countermeasures. “People don’t understand COVID-19 simply because it is invisible. Our goal is for our HPC results to help people to be more informed, especially government planners who will establish guidelines for a post-COVID society,” he added.

As the pandemic continues to rapidly unfold, it may be hard to see the light at the end of the tunnel. But the coronavirus may have finally met its match in the form of supercomputers like ASPIRE1, Fugaku and Nurion. With the world’s best minds and fastest supercomputers joining forces to attack COVID-19 on all fronts, there’s still hope yet. ■

CALLING ALL SINGAPORE-BASED SCIENTISTS!

In light of the coronavirus outbreak, the National Supercomputing Centre (NSCC) Singapore is opening its doors to all local researchers working on COVID-19-related projects, including both existing NSCC users and non-users. Applications will be reviewed by a selection committee chaired by Dr. Kenneth Ban, the NSCC Director for Health & Biomedical Sciences Initiatives. Resources will be provisioned to approved applications within two weeks from the submission date and will be valid for one year from the date of allocation.

RESOURCES AVAILABLE AT THE NSCC:

Resource	Specifications
ASPIRE 1	1,288 CPU and 128 accelerator nodes with NVIDIA K40 GPUs and 31,392 cores
AI system	Six units of eight-GPU NVIDIA DGX-1 Deep Learning System with V100 cards
High performance storage	13 petabytes

Hurry, applications are open only until September 23, 2020! For more details, visit help.nscg.sg/anticovid or contact NSCC at projects-admin@nscg.sg or bizdev@nscg.sg

SCAsia

Supercomputing 2021

SAVE THE DATE BECAUSE IT'S ON!

SINGAPORE 1-4 MARCH 2021

Suntec Singapore Convention
& Exhibition Centre

MAPPING THE BRAIN'S MYSTERIES



Harnessing HPC to understand the brain

Scientists across Asia are harnessing the power of synchrotrons and supercomputers to get a comprehensive understanding of the brain.

By **Tim Hornyak**

Illustrations by Lam Oi Keat/Supercomputing Asia

The most complex structure in the known universe is a paradox. It's right between our ears, but much of it remains beyond our understanding. The human brain contains over 86 billion nerve cells, or neurons, and some 100 trillion connections among them. Add to that neural pathways, glial cells and messenger molecules called neurotransmitters, and you've got a very intricate and powerful organ, one capable of probing the structure of the cosmos and laughing at silly jokes.

Scientists in Asia are now bringing the world's most sophisticated computer technology to bear on this eternal mystery by comprehensively mapping the human brain for the first time—by as soon as 2024. This three-dimensional map, known as a connectome, could have a huge impact on everything from Alzheimer's disease to artificial intelligence. From studies of monkey brains, to mapping the human brain and even simulating human neural activity on supercomputers, this effort could have profound consequences for both brain science and some of our fundamental questions about what it means to be human.

AN ASIAN CONNECTOME CONSORTIUM

In January 2020, the National University of Singapore played host to the inauguration of a bold collaborative effort to create a brain map. Synchrotron for Neuroscience—an Asia Pacific Strategic Enterprise (SYNAPSE) brings together researchers across the region to image the brain at the sub-cellular level, specifically 0.3 micrometer resolution, at a speed of one cubic millimeter per minute. Since capturing that level of detail would take years at a single facility, the work is being divided among several centers in the region while using automation and joint management of data to accelerate the project.

"Mapping the human brain at a resolution sufficient to chart the connections is a historic mission for science and technology," said Low Chian-Ming, an associate professor at the National University of Singapore and a founding member of the initiative. "SYNAPSE will also generate technological breakthroughs in imaging, computation and artificial intelligence."

More than 1,000 researchers are expected to participate in the undertaking, hailing from institutions such as Academia Sinica in Taiwan; Pohang Light Source-II in South Korea; Spring-8 in Japan; Singapore Synchrotron Light Source at the National University of Singapore; Shanghai Synchrotron Radiation Facility in China; and Australian Synchrotron.



The main imaging technology they will use is X-rays from synchrotrons, giant rings that can accelerate electrons close to the speed of light. When moving electrons change direction, they release high-energy X-rays, which can be shunted through solid objects to image their internal structures. "Synchrotron-based imaging is essential to achieving meaningful connectome mapping of a large animal brain because it is the only technique that combines three-dimensional sub-cellular resolution with high imaging speed to collect data for an entire brain within a reasonable time," said Professor Hwu Yeu-Kuang, leader of the Taiwanese team at Academia Sinica.

Low and colleagues are confident of their approach because it builds on previous work under Accelerated X-ray Observation of Neurons (AXON), a set of X-ray

imaging technologies perfected by an international scientific collaboration which is now the core of the SYNAPSE project. This technique demonstrated the speed and resolution possible on the fruit fly brain and is on its way to being tested on larger animal brains. The researchers also managed to partially map a mouse brain, according to a study in the *Chinese Journal of Physics*.

"Our first data set for the fruit fly brain, published in a study led by Professor Chiang Ann-Shyn at the National Tsing Hua University, proved that even results for small brains and partial information can revolutionize neurobiology research," Hwu added.

SYNAPSE will incorporate petascale high performance computing facilities at the National Supercomputing Centre (NSCC) Singapore, which is a sponsor of the project. SYNAPSE is expected to produce a huge volume of data that will be in excess of one billion gigabytes, or one exabyte, which is the equivalent of a million one-terabyte drives. The NSCC will serve as a data hub for the processed 3D mapping data, linking processing facilities like Fugaku at the RIKEN Center for Computational Science, Taiwan at Taiwan's National Center for High-Performance Computing, and systems at the Daegu Gyeongbuk Institute of Science and Technology, South Korea, via a high-speed 100 Gbps network run by the Singapore Advanced Research and Education Network (SingAREN).



High resolution tomographically reconstructed image of a *Drosophila melanogaster* fly head showing the nervous system (green), muscles (orange), cuticles (gray) and compound eye. The reconstruction was performed from 600 projection images taken at equal intervals within 180°.

Credit: Reproduced with permission from Chin et al (2020) A synchrotron X-ray imaging strategy to map large animal brains.

“Our task is to provide neuroscientists with something like a Google Map of a complete functional brain, where they can view different levels of information as they zoom in.”

“We specifically tested the parallel involvement of different synchrotrons and their calibration, so that data can be simultaneously analyzed and processed with uniform high performance computing platforms to be shared among partners,”

said Low. “Our task is to provide neuroscientists with something like a Google Map of a complete functional brain, where they can view different levels of information as they zoom in, balancing the need for detailed information along with retrieval speed. It will also serve as a reference to chart the functional readings obtained through fMRI, electrophysiology, EEG and other techniques, building up to a structure-function brain map.”

Other imaging techniques that the SYNAPSE researchers are using include infrared spectromicroscopy, super-resolution visible light 3D microscopy and cryo-electron tomography. The information on specific brain regions obtained from these complementary techniques would then be used to annotate the brain map with a large variety of functional and high-resolution information related to the neural network.

In being able to shed this much light on the structures of the brain, SYNAPSE may give rise to something that’s long been a staple of science fiction: simulating a human brain on a computer. Since it will be based on neural network data from real brain maps, it’s expected to generate more understanding about brain functions, complementing initiatives such as the BRAIN Initiative in the US and the Human Brain Project in Europe, according to Low.

“No individual country would have the resources to finish this project in a reasonable amount of time. Parallel data collection is essential and can only be realized by an international project,” Hwu said. “The same is true for processing and storing the huge amount of data, and for further developments of the imaging and data processing techniques.”

In the same spirit, the results will be shared by all partners and open to everyone, with the help of cloud

computing and high performance computing. “As the raw data size is in exabytes, each country partner will store the raw data but deposit the reconstructed core data at NSCC Singapore,” Low explained.

In the first phase that is expected to be completed by 2024, SYNAPSE aims to map one human brain. “However, the main goal of our collaboration is to understand how the human brain functions from the complete mapping of the neural network. This could lead to effective therapies for brain diseases, a global problem with a huge social impact,” Low said. “This mission may take more than a decade.”

CHINA’S LONG PLAY

SYNAPSE follows a number of other large-scale brain computing endeavors in East Asia. One is the China Brain Project (CBP; also known as the China Brain Initiative), an enormous initiative stretching to 2030 with multiple aims: expanding basic research on cognitive functions; applied research into the diagnosis and treatment of brain disorders; and research into computing methods inspired by the human brain.

The CBP is intended to rival similar efforts in the West, and promoters have said that compared to other initiatives, the CBP is prioritizing brain disorders and artificial intelligence as its near-term goals. Part of that is due to the fact that China, with its enormous and aging population, is saddled by large numbers of people with neurological and cognitive disorders.

Professor Poo Mu-ming, leader of the CBP and director of the Institute of Neuroscience of the Chinese Academy of Sciences, said the CBP has not officially started yet, and referred *Supercomputing Asia* to a 2016

article he co-authored in the journal *Neuron* which stated that the CBP is part of a 15-year plan.

“The China Brain Project covers both basic research on neural mechanisms underlying cognition and translational research for the diagnosis and intervention of brain diseases as well as for brain-inspired intelligence technology,” Poo and colleagues wrote.

While there have been few details of this grand national undertaking, a number of new research centers have opened. The Chinese Institute for Brain Research, Beijing, was inaugurated in 2018 as a main pillar of the CBP. It has been engaged in recruiting 50 internationally selected principal investigators and more than 1,000 other researchers as well as building 11 core facilities with some 30,000 animal cages. A sister center, the Shanghai Research Center for Brain Science and Brain-Inspired Intelligence, was also launched in 2018.

The CBP’s emphasis on animal research comes after a team led by Poo produced two genetically identical long-tailed macaques, a world first for primates, using somatic cell nuclear transfer, the technique that created Dolly the sheep. The study, reported in the journal *Cell*, raised hopes that the cloning know-how could be used with gene-editing techniques such as CRISPR-Cas9 to recreate monkey models of human brain disorders such as Parkinson’s disease.

“This paper really marks the beginning of a new era for biomedical research,” Professor Xiong Zhi-Qi, a neuroscientist at the Chinese Academy of Sciences Institute of Neuroscience, was quoted as saying by *Nature News*. Similar work has broadened the field: Chinese scientists have even created transgenic macaques with extra copies of a human gene that might have a role in human intelligence. They reported that the enhanced monkeys outperformed their peers on a memory test. As if alluding to science-fiction scenarios from *Planet of the Apes*, Western scientists have questioned the ethics behind tinkering with monkey genes and intelligence.

SUPERCOMPUTING BRAINS IN JAPAN


Unlike the CBP which focuses on macaques, Japan’s Brain Mapping by Integrated Neurotechnologies for Disease Studies (Brain/MINDS) looks at another non-human primate: marmosets. Brain/MINDS was launched in 2014 and is now in its second phase. With the state-backed RIKEN Center for Brain Science (RIKEN CBS) playing a coordinating role, Brain/MINDS brings together universities and research institutes from throughout the nation.

The project’s goal is to elucidate the neural circuits underlying higher brain functions using unique experimental models. A particular focus is neural circuits that are responsible for neurological and psychiatric disorders. As member researchers wrote in *Philosophical Transactions of the Royal Society B: Biological Sciences*, the objectives fall into three main areas: functional mapping of the marmoset brain, developing innovative neurotechnologies for brain mapping and actual brain mapping, along with clinical research.


Recent scientific reports from the initiative include the role of auditory signal processing in schizophrenia, neuroinflammation in mouse models of Alzheimer’s disease, and a comparative study of how common marmosets and Japanese macaques react to human behavior. RIKEN’s Professor Hideyuki Okano, dean of Keio University’s Graduate School of Medicine, has been producing genetically modified marmosets to learn more about Rett Syndrome and Parkinson’s disease.

Tomographically reconstructed image of a portion of a mouse brain from 600 projection images. The different colors correspond to different neuron clusters, each formed by interconnected cells.

Credit: Reproduced with permission from Chin *et al* (2020) A synchrotron X-ray imaging strategy to map large animal brains.



The human brain contains approximately 100 billion neurons



The cerebral cortex occupies 82% of the brain mass, but only 19% of all neurons

The cerebellum occupies only 10% of the brain mass, but contains 80% of all neurons

Therefore, simulating the cerebellum is akin to simulating 80 percent of the whole brain

“The marmoset brain is essentially a simpler primate brain to study and its common features to other primates make it easier to analyze,” said Professor Alexander Woodward, leader of the Connectome Analysis Unit at RIKEN CBS, where the goal is to chart and analyze brain connections. “The developmental period is much shorter than that of humans. Furthermore, the brain does not have the gyrification (wrinkled appearance) that larger primates or the human brain have, making it easier to carry out certain invasive experiments.”

A key part of these experiments is tracer injections. As Woodward explains, a viral tracer injected into the brain infects neurons, causing them to express a fluorescent protein that travels down the axon to the terminal sites. By later examining the brain *ex vivo*, researchers can see the connections that neurons at the injection site make with neurons in other brain regions.

Brain/MINDS scientists are also using MRI technology known as diffusion weighted imaging (dMRI) to obtain structural data. This data can provide insights into the overall connectivity pathways and the structural pathologies where there is damage to the brain connections. Members of Woodward’s team also developed a 3D marmoset brain atlas which describes all of the brain regions and this is being used to calculate and summarize the brain connectivity patterns from the data.

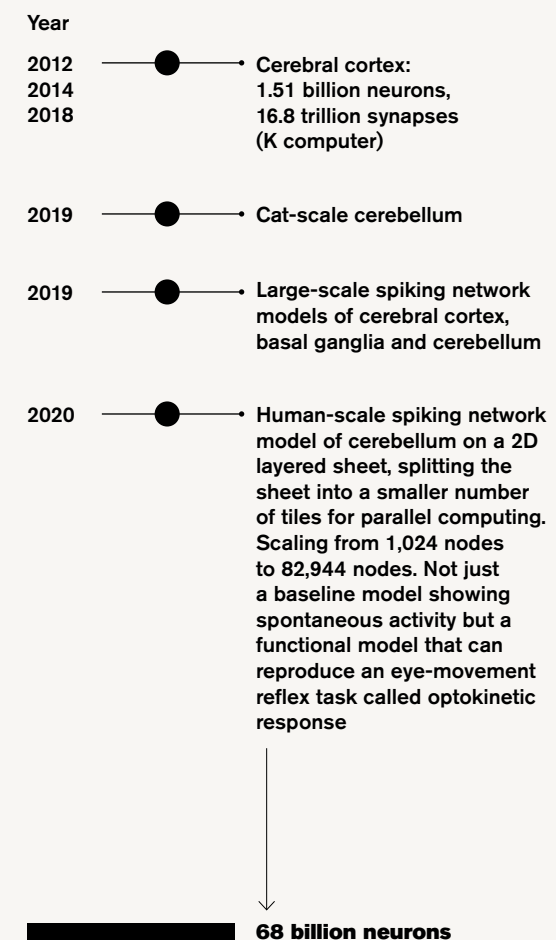
PUSHING THE FRONTIER WITH FUGAKU

Another major RIKEN project related to the Brain/MINDS endeavor is the construction of a next-generation supercomputer named Fugaku (see our infographic on p. 38). An alternate name for Mount Fuji, Japan’s tallest peak, Fugaku is designed to have an ARM architecture and some 150,000 CPUs. It would be at least 40 to 100 times more powerful than its predecessor, the K computer, which TOP500 ranked as the world’s fastest supercomputer in 2011. Under construction by Fujitsu, Fugaku could become the world’s first computer with exaflop speed, or 10¹⁸ floating-point operations per second, if it enters full operations around 2021 as planned and doesn’t lose out to rivals in China, the US and the European Union.

Fugaku will be used for everything from drug discovery to simulating earthquakes and tsunamis. Neuroscientists in Japan, though, are eager to use it for brain simulations. They believe Fugaku will be the first computer with sufficient resources to simulate a human brain—it will recreate the same number of neurons as the human brain as well as the connections between them. The feat will be astonishing given that only seven years ago, it took the K computer 40 minutes, using 1.73 billion virtual nerve cells connected by 10.4 trillion

“Our goal is to realize whole-brain simulation using the Fugaku computer for an understanding of the mechanisms of brain function and disease.”

Previous attempts at simulating the brain:



synapses, to simulate only one second of neuronal activity in a human brain. Japanese researchers have already begun simulations on part of Fugaku and, with the full system, they plan to test the computational performance of a brain model consisting of the cortex, cerebellum, thalamus and basal ganglia.

“Our goal is to realize whole-brain simulation using the Fugaku computer for an understanding of the mechanisms of brain function and disease,” said Dr. Jun Igarashi of the High Performance Artificial Intelligence Systems Research Team at RIKEN’s Center for Computational Science. “As a first step, we plan to perform whole-brain simulation of mouse, marmoset and human, to understand the interactions between brain areas in processing information for movement and behavior.”



Using machines like Fugaku, scientists may be able to better understand the brain’s fundamental information-processing ability as well as activity, such as how movement is coordinated over a span of seconds to minutes. But will researchers actually be able to simulate a mind on a computer? That remains to be seen.

“No one knows whether whole-brain simulation can reproduce human intelligence until it is realized,” Igarashi said. “The current problem is that there is no consensus on which levels of description of the neuron are sufficient for reproducing intelligence.”

Bolstered by grand initiatives like these projects across Asia, scientists hope to one day gain a complete mechanistic understanding of the universe’s most complex object. But that doesn’t mean our brains can be reduced to biological computers. If anything, it’s the part of us that can’t be simulated—the ghost in the machine—that makes us human.

GETTING A HEADSTART ON HPC

High performance computing at institutes of higher learning



As high performance computing grows in relevance in various industries, Singapore Polytechnic is prepping its students for the future from the comfort of their classrooms.

By **Kamila Navarro**

First introduced in 1954 to meet the country's need for skilled labor, polytechnics have played a pivotal role in Singapore's successful industrialization. But the world has profoundly changed since then, particularly following the invention of the integrated circuit in 1959. Today, powerful supercomputers with tens of thousands of chips are nothing out of the ordinary. With governments and industries alike increasingly facing complex problems answerable only with high performance computing (HPC), polytechnics must quickly adapt to equip their students with the skills to meet the latest industry needs.

To help their students stay ahead, Singapore Polytechnic (SP) has launched a full-time Diploma in Data Analytics and Artificial Intelligence (DAAA)—the first of its kind among the five polytechnics in Singapore.

But DAAA offered by the School of Computing (SoC) isn't your run-of-the-mill data science/AI program. HPC, for example, features quite prominently in DAAA's Practical AI module, enabling students to gain first-hand experience in a concept that had hitherto been thought inaccessible.

BUILDING AN HPC-READY WORKFORCE

"Many industries, especially consumer companies, are deploying machine learning and deep learning for their production lines," explained Dr. Edna Chan, center director of SP's Data Science and Analytics Centre, in an interview with *Supercomputing Asia*. "For such applications, especially those that are computationally expensive, they won't be using normal servers and personal computers."

Industry giants like General Electric (GE) and French Big Oil company Total have already begun investing in their own in-house supercomputing assets to improve their products and processes. GE, for instance, has used HPC to improve the design of their wind turbines, making them more aerodynamic and energy-efficient. Meanwhile, Total's Pangea III—also known as the world's largest industrial-use supercomputer—has been used to improve the detection accuracy of underground hydrocarbons.

Considering the close relationships between SP and its industry partners, Chan and her colleagues decided that it was high time for their students to gain the HPC skills desired by industry. "It's important for SP students to understand how HPC can be used in such situations and applications and be ready to make use of HPC when they enter the workforce," she added.

This necessity for HPC skills was taken into consideration by SP as they developed the DAAA program. Hence, in the Practical AI module, DAAA students will specifically learn to perform the batch management and checkpointing of large-scale AI training jobs, allowing them to truly put HPC to the test. However, other students, like those from the Diploma in Computer Engineering (DCPE) will also get to join in the fun. "As HPC hardware is typically housed in data centers, DCPE students will also learn to manage these centers," said Chan.



HIGH-POWERED HELP

Moving beyond industry, it's worth noting that HPC is also no longer limited to advanced academic research laboratories. HPC has popped up in more mainstream applications, ranging from product design to driverless vehicles. Interestingly, some of these applications are currently being explored by SP's various schools to address concerns on campus.

For example, a team from the School of Electrical & Electronic Engineering (SEEE) is working on a driverless electric vehicle that contains a bevy of sensors, including cameras, 5G and a light detection and ranging (LIDAR) system. Given the volume and complexity of the vehicle's data, the team intends to use HPC to speed up

Singapore Polytechnic students from the School of Electrical & Electronic Engineering with their HPC-powered autonomous electric vehicle.

computationally intensive processes like training the vehicle to recognize objects in its path. Once ready, the vehicle will be deployed around the SP campus to transport passengers and serve as an autonomous campus control platform. Given the current COVID-19 situation, the driverless electric vehicle could also double as a means to monitor and remind students, staff and visitors to adhere to precautionary measures like maintaining social distancing and wearing face masks.

SEEE has also developed a campus surveillance and building inspection system that utilizes 5G to analyze drone video footage in real-time. Just like the previous example, HPC will be used to train the complex images

required by the system. With the help of HPC resources, Chan notes that SP can now take on more demanding projects and dream bigger, making the long-awaited move from mere experimentation to the level of minimum viable product.

AHEAD OF THE PACK

The School of Computing is planning to build two AI and Analytics Colabs equipped with 50 deep learning graphics processor units (GPU)-powered workstations and servers. Meanwhile, SEEE's machine learning and AI laboratory will have 25 NVIDIA RTX GPU-powered workstations for training and modeling deep neural networks—an evolution of machine learning where brain-inspired artificial neural networks learn from large amounts of data. As indicated by its name, GPUs were originally intended to accelerate graphics tasks like rendering images, making them incredibly popular among avid video gamers. GPUs, however, have since found a second life in HPC due to their ability to carry out massive parallel processing.

While SP's current assortment of equipment isn't massive enough to be considered an HPC cluster, Chan shares that it should still allow students to appreciate the possibilities offered by HPC, especially when it comes to data-intensive modeling. "Students can leverage HPC resources graciously provided by the National Supercomputing Centre when far more intensive computation is required," she said. While SP may be the oldest polytechnic in Singapore, its students are nonetheless kept up to date with a rapidly moving industry and even given a head start on HPC. ▣

INTEL SCALES NEUROMORPHIC SYSTEM TO 100 MILLION NEURONS

Inching closer towards its goal of mimicking the human brain, Intel revealed last March its largest and most powerful neuromorphic research system to date. Called Pohoiki Springs, the cloud-based system has a computational capacity of 100 million neurons—roughly the same number as in the brain of a hamster.

Pohoiki Springs integrates 768 Loihi neuromorphic research chips into a five-rack unit. As suggested by its name, neuromorphic chips are modeled after the brain's architecture. Neurons and synapses are simulated using artificial neural networks, allowing computers to learn through observation rather than programming. Incredibly, neuromorphic computing systems are also more energy-efficient than traditional hardware.

Currently, Pohoiki Springs and Intel's other neuromorphic systems are still in the research phase and not meant to replace the conventional computing systems we see today. But given the breadth of its scale, Pohoiki Springs can be used to develop advanced brain-inspired algorithms. Examples of algorithms being developed for its Loihi chips include solving games like Sudoku and matching patterns for face recognition.

Intel's previous efforts had produced chips on the smaller end of the scale. Kapoho Bay, for instance, only has two Loihi chips and the equivalent of 262,000 neurons. In contrast to Pohoiki Springs, Kapoho Bay is intended for edge algorithms.



Arvind Krishna
Chief executive officer, IBM

Jim Whitehurst
President, IBM

IBM UNDERGOES LEADERSHIP RENEWAL

After eight years at the helm, IBM's CEO and president Virginia 'Ginni' Rometty has stepped down. Taking over as the technology conglomerate's new leader is Arvind Krishna, who currently serves as IBM's senior vice president for cloud and cognitive software. Red Hat CEO Jim Whitehurst will take over as president.

A well-known business leader, Rometty took over from Sam Palmisano in 2012. During her tenure, she made cloud computing a priority for IBM and also oversaw the integration of PwC with IBM in the early 2000s. She began her career at IBM in 1981 and held roles in sales, marketing and strategy. Rometty will continue to serve as executive chairman until the end of the year.

Krishna, who joined IBM in 1990, will be the tenth CEO of the organization. It was Krishna who drove the massive US\$34 million acquisition of open source software provider Red Hat, making history as IBM's largest acquisition to date. Whitehurst, on the other hand, is known for leading Red Hat to become the first US\$1 billion company of its kind.


SAMSUNG ANNOUNCES LAUNCH OF POWERFUL NEW MEMORY CHIP

Imagine being able to download 82 full HD movies in under 20 minutes. That's the kind of speed that Samsung's latest memory chip boasts of.

Known as the 'Flashbolt,' it is the third-generation high bandwidth memory (HBM2E) technology to be launched by the South Korean technology giant. Flashbolt could find takers in industries that rely on high performance computing systems like supercomputers and artificial intelligence. Set to be launched later this year, the chip is 1.3 times faster and has double the memory capacity of its predecessor, Aquabolt.

With 16 GB of memory and a data transfer speed of 3.2 GB per second, the Flashbolt could be a shot in the arm for the next generation of supercomputers that are hungry for both speed and capacity. When overclocked, these memory chips can offer up to 4.2 Gbps data transfer rates per pin and a memory bandwidth of 538 Gbps. For comparison, the Aquabolt offered 2.4 Gbps bandwidth per pin and 307.2 Gbps bandwidth per stack.

Set to go into mass production within the year, these DRAM chips will likely be used in upcoming high-end GPUs from chip makers like AMD and NVIDIA.



ATOS ACQUIRES CLOUD CONSULTING FIRM MAVEN WAVE

Information technology consulting company Atos has completed its acquisition of Maven Wave, a Chicago-based cloud and technology consulting firm with strengths in building cloud-based digital solutions, machine learning and data analytics. While the deal was announced in late 2019, Atos completed the acquisition in February of this year, strengthening its position in the North American market.

The acquisition also reinforces Atos' position as a proven leader in the cloud computing industry. In 2019, it was named the Google Cloud Breakthrough Partner of the year and adds 330 Maven Wave employees into its fold. The Paris-headquartered company currently holds ten Google specializations and has retained the title of Google Cloud North America Services Partner of the Year for two years in row.

"Together, Maven Wave and Atos create the strongest Google Cloud services portfolio offered anywhere, providing customers proven expertise and knowledge in executing their digital transformation and delivering outstanding experiences to their customers," said Maven Wave founders Briar Farrar, Jason Lee and Jeff Lee in a company statement.

Super Snapshot

ON TOP OF THE WORLD

Named after Mount Fuji—one of Japan's three holy mountains and its tallest peak—the Fugaku supercomputer will succeed the decommissioned K as Japan's next-generation, flagship supercomputer. And like Mount Fuji, Fugaku is expected to be a towering figure in the supercomputing world. Installed at the RIKEN Centre for Computational Science in Kobe, Japan, Fugaku is currently the fastest supercomputer in the world.

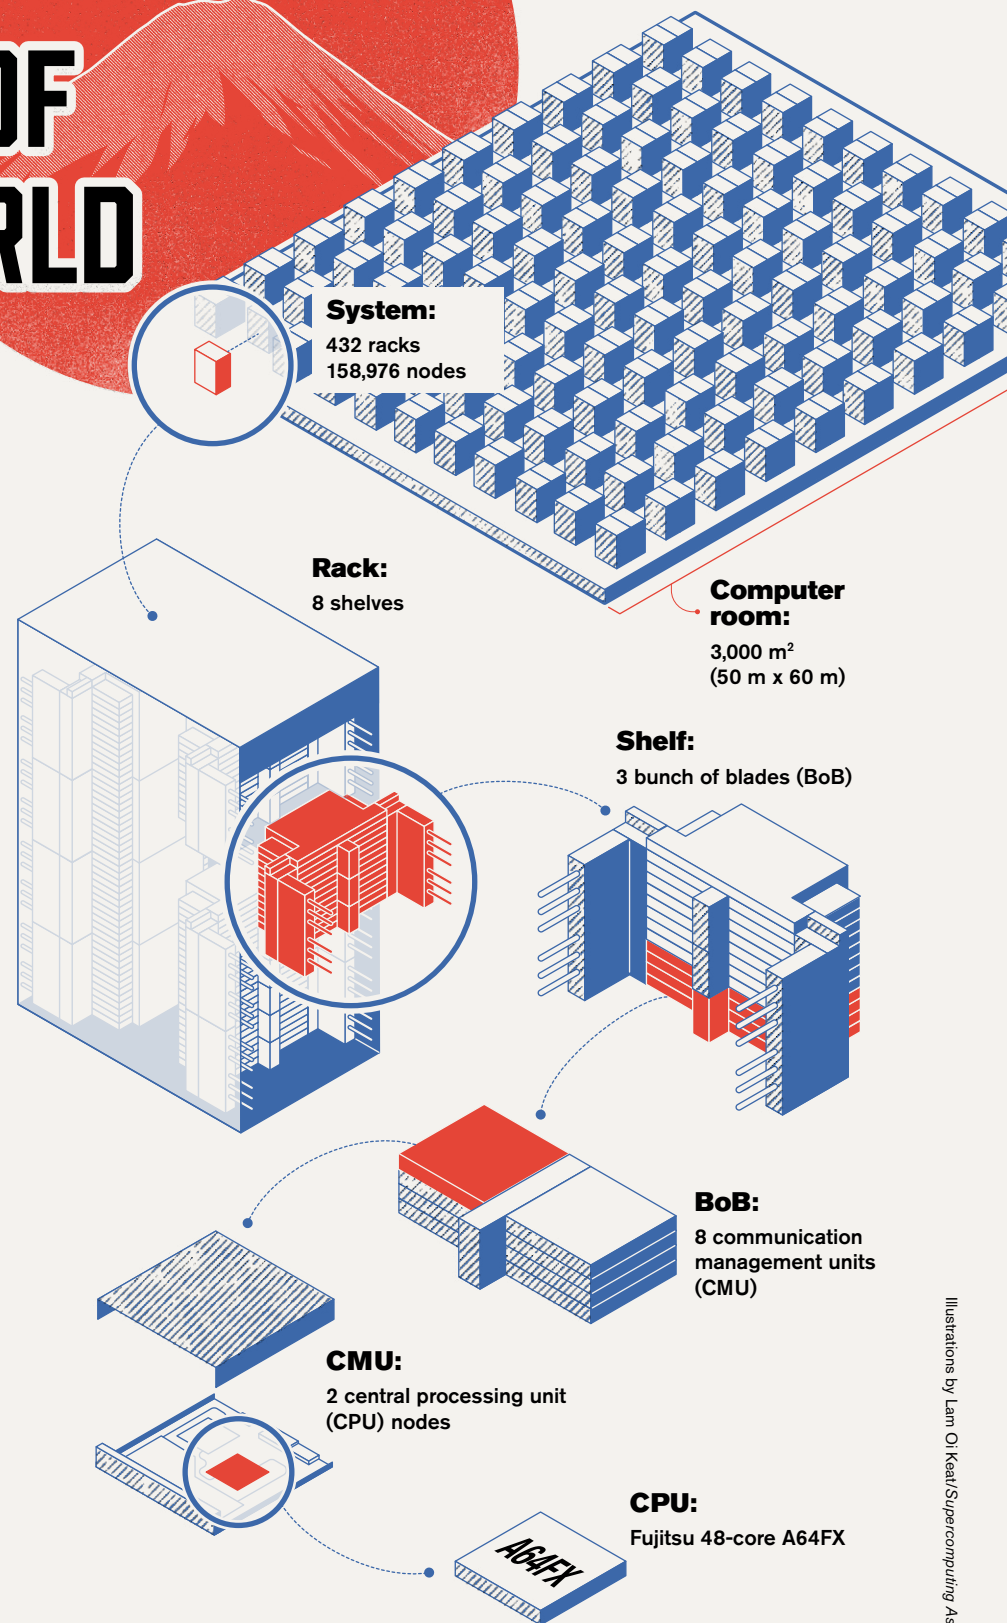
Fugaku leads across four HPC ranking tables:

- #1 TOP500: 415 petaFLOPS
- #1 HPCG: 13.4 petaFLOPS
- #1 HPL-AI: 1.42 exaFLOPS
- #1 Graph500: 70.98 teraTEPS

Priority areas:

- 1 Drug discovery
- 2 Personalized & preventive medicine
- 3 Tsunami & earthquake prediction
- 4 Meteorological prediction
- 5 Efficient energy creation, conversion & use
- 6 Clean energy
- 7 New materials
- 8 Innovative design & production processes
- 9 Fundamental laws & evolution of the universe

Look out for an interview with the architect of Fugaku, Professor Satoshi Matsuoka, in the next issue of Supercomputing Asia!



Illustrations by Lam Oi Keat/Supercomputing Asia

COMPETE & COMPUTE

The HPC-AI Advisory Council
2020 APAC HPC-AI Competition Is On!
30 University Teams, 9 Regions, One Mission: Mastery in HPC and AI.

The teams will face some of the most difficult real-life High-Performance Compute challenges, like COVID-19 medical research and more.

Mark your calendar: the Official LIVE award ceremony
SupercomputingAsia 2021 Conference March 2021, Singapore.



SUPERCOMPUTING FOR SINGAPORE

NSCC Singapore was established in 2015 and manages Singapore's first national petascale facility with available high performance computing (HPC) resources. As a National Research Infrastructure funded by the National Research Foundation (NRF), we support the HPC research needs of the public and private sectors, including research institutes, institutes of higher learning (IHLs), government agencies and companies. With the support of its stakeholders, NSCC catalyses national research and development initiatives, attracts industrial research collaborations and enhances Singapore's research capabilities.

Find out more about our resources and our recent Project Calls including access to supercomputing power for COVID-19 research work and for Educational HPC activities

📍 1 Fusionopolis Way, Connexis South Tower, #17-01, Singapore 138632

Want to tap on supercomputing power?
See what project calls are ongoing at
<https://help.nsc.sg/>

📘 NSCCSG