ASIAN SCIENTIST

Issue 03
January 2018

SUPERCOMPUTING ASIA

SUPERCOMPUTERS TAKE TO THE SEA

A NEW AGE OF MARINE EXPLORATION

+

RIDING THE HPC WAVE

THE VIEW FROM THE TOP500

COMPUTING IN THE COSMOS

MPAS AWARDS 2017
Technology Trade Media of the Year
Science Trade Media of the Year

THE 10TH IDA INTERNATIONAL DESIGN AWARDS
Silver Award

MCI (P) 166/04/2017

CONTENTS

Issue 03
Jan 2018

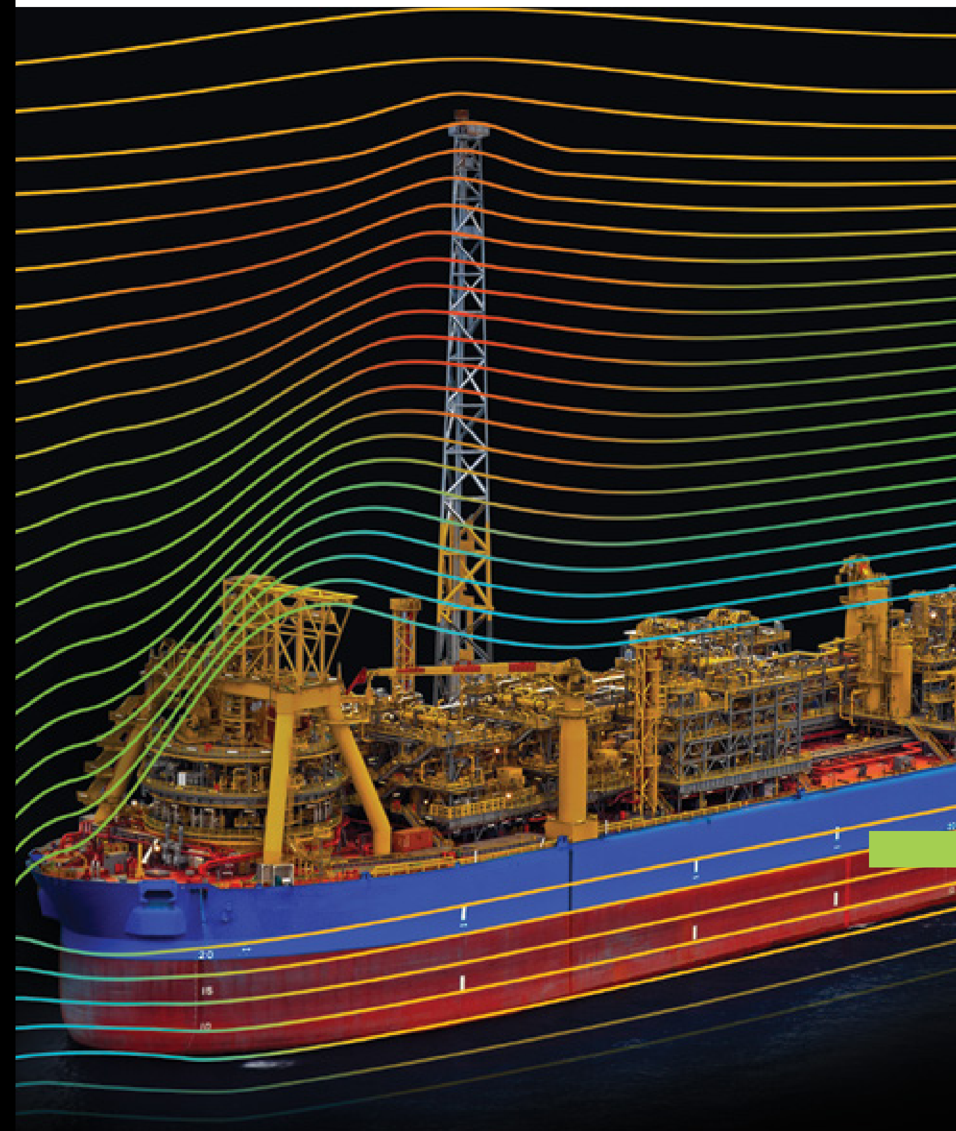

EDITOR'S NOTE

More than 70 percent of the Earth's surface is covered by water, the vast majority of which can be found in the oceans. And yet, the seas and their depths remain relatively unknown and largely unexplored.

Supercomputers are slowly but surely beginning to chip away at our ignorance of the ocean. This issue, we survey how simulations are ushering in a new age of marine exploration (*Supercomputers Take to the Sea*, p. 16) and five ways that supercomputers are already being put to use in the oil and gas industry (*Riding the HPC Wave*, p. 10).

From the ocean depths, we turn our attention to outer space in our interview with Hewlett Packard Enterprise's Dr. Goh Eng Lim (*Computing in the Cosmos*, p. 26). Last but not least, check out our interview with Professor Jack Dongarra as he shares his unique perspective on the progress of supercomputing drawn from 25 years of maintaining the TOP500 list (*The View from the TOP500*, p. 22).

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

/asianscientist @asianscientist
Asian Scientist Magazine

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

SENIOR EDITOR
Dr. Sim Shuzhen

ASSISTANT EDITOR
Dr. Jeremy Chan

CONTRIBUTORS
Jerry Lim
Jon Lau
Cindy Lim

SENIOR DESIGNER
Lam Oi Keat

JUNIOR DESIGNER
Allan Chan

SALES & MARKETING
Samantha Yeap

PUBLISHED BY
Wildtype Media Group Pte Ltd

DISTRIBUTED BY
Pansing Distribution Pte Ltd

HEAD OFFICE
Wildtype Media Group Pte Ltd
71 Ayer Rajah Crescent #07-03
Singapore 139951
Tel: 65-6909-2279
Website: www.asianscientist.com/subscribe

WILDTYPE MEDIA

To order reprints or e-prints, or request permission to republish Supercomputing Asia content, please contact admin@asianscientst.com

Supercomputing Asia, MCI (P) 166/04/2017, 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.

SINGAPORE'S First Petascale Supercomputer

ASPIRE 1

(Advanced Supercomputer for Petascale Innovation Research & Enterprise)

Gain Faster & Better Insights into product design, leading to robust product performance

Reduce Total Cost of Ownership by using HPC resources on an on-demand basis to provide agility to run organisations in a rapidly changing business environment

Decrease Time-to-Solutions for optimised corporate IP research and shorten time-to-market to accelerate commercialisation of products

Enable High Fidelity modeling and simulation and large-scale data-driven researches and application for commercialisation

ASPIRE 1 USERS

Mr Paul Jones, Chief Commercial Officer, Global Gene Corp
"The compute power is huge and has allowed us to have the most optimal turnaround time for analytical work performed on advanced bioinformatics analysis of large genomic datasets. We certainly have several work plans in the pipeline to further explore our collaboration with NSCC!"

Arsen Batagov, Bioinformatics Director, Vishuo Biomedical
"NSCC has provided us with the infrastructure to run the analysis on hundreds of CPU cores, thus reducing the execution time from days to hours. It has also given us the advantage of being able to deliver the analysed results to our customers faster, thus making us a strong competitor within global markets for bioinformatics services."

Students, National Junior College
"The ASPIRE1 supercomputing system had to be used as the parameters of the calculations are too 'large' for "normal" machines to compute within a short time frame. ASPIRE 1 allowed us to circumvent these challenges and conveniently by simply logging into the supercomputer system using our laptops."

SUBSCRIPTION PLANS AT A GLANCE (Industry)

	BASIC	SILVER	GOLD
No. of user accounts	1	3	5
Maximum active jobs per user	100	100	100
Maximum active jobs per tier	100	300	500
Maximum CPU-cores per job	400	600	800
Maximum wall time per job	24 hours	24 hours	24 hours
Access to long queue (24hrs – 240hrs)	No	Yes	Yes
Job queueing priority	Standard	Medium	High
Indicative maximum CPU hours	5,000	20,000	50,000
Storage quota per tier (GB)	50	1,000	3,000
Package Price	**S$ 500**	**S$ 2,000**	**S$ 5,000**

+ Add-on Services	Price / Hour
CPU (per core)	S$ 0.1070
GPU (per node)	S$ 1.1860
Storage (GB)	S$ 0.0323 (minimum of 50 GB)

All prices are subject to 7% goods and services tax.

SG NSCC SINGAPORE
National Supercomputing Centre

For business opportunities, please contact the NSCC Business Development team:
bizdev@nscc.sg or +65 6714 9450

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

www.nscc.sg

Digital Dispatch

CHINA BAGS SECOND GORDON BELL PRIZE

For the second time in the history of the highly coveted Gordon Bell Prize, a team from China has won the award, receiving their commendation at the International Conference for High Performance Computing, Networking, Storage and Analysis (SC17) on November 16, 2017.

The winning simulation was performed on Sunway TaihuLight, currently the world's most powerful supercomputer, according to the fiftieth TOP500 list. Sunway TaihuLight was also used for the project that won last year's prize.

This year's winners were recognized for their simulation of the 1976 Tangshan earthquake, one of the most devastating earthquakes of the 20th century. The team developed software that was able to process 18.9 petaFLOPS of data from the entire area of the quake, a surface diameter of 320 km by 312 km stretching down to 40 km below the earth's surface.

Such detailed, high-frequency 3D simulations could be used to guide building standards for areas with known seismic activity. In addition, the on-the-fly compression scheme they developed could be applied towards other challenges in exascale computing, the researchers said.

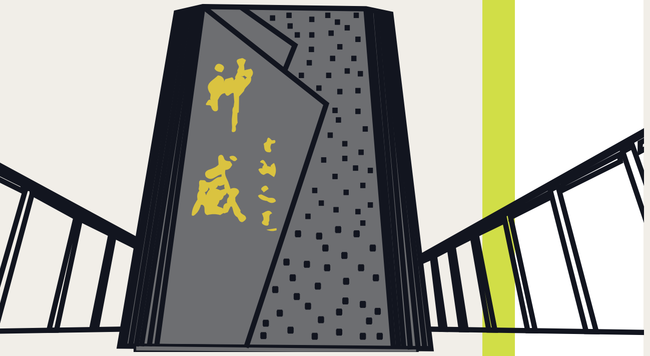

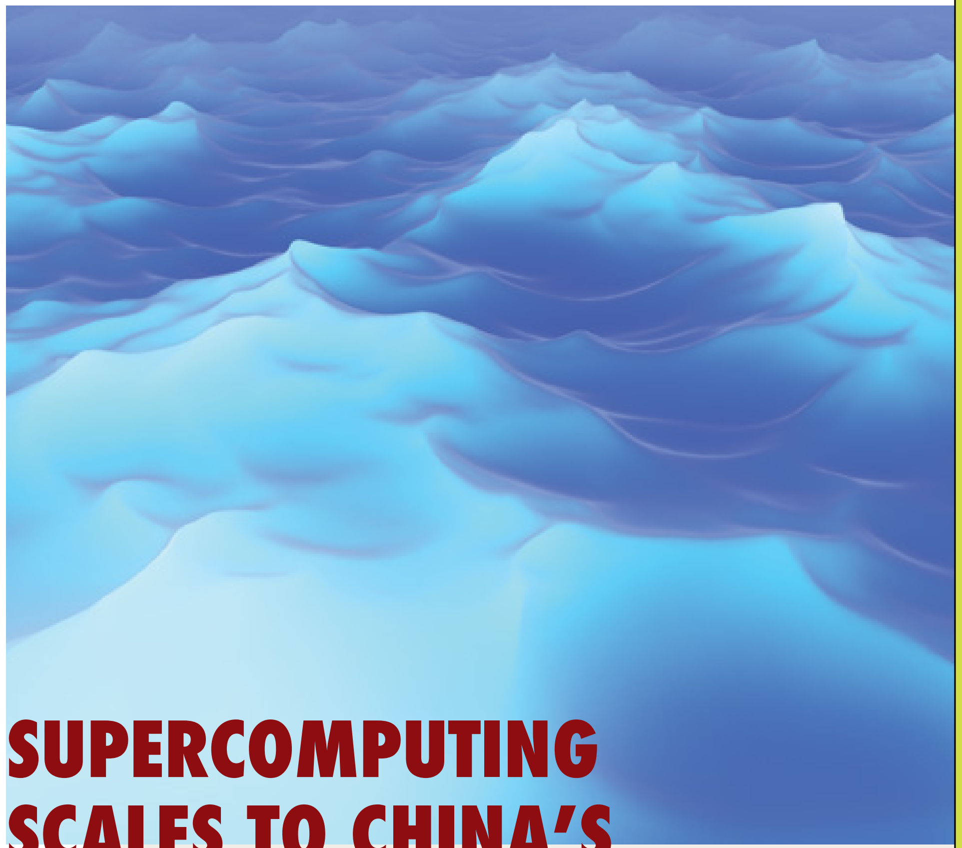

SUPERCOMPUTING SCALES TO CHINA'S MARINE AMBITIONS

The third of China's three exascale prototype supercomputers will be used for marine research, according to a report by the *South China Morning Post*.

In a bid to develop an exascale supercomputer for China by 2020, three different groups are currently working on prototypes. The first group, led by Dawning Information Industry Co Ltd (Sugon), announced its plans in late 2016.

In January 2017, the second group—the National University of Defense Technology—announced that their Tianhe-3 prototype would be ready by the end of the year.

The third contender, being developed by the team behind the world's fastest supercomputer, Sunway TaihuLight, is expected to launch its prototype by June 2018. In addition to space science and medical applications, it will be used by the Institute of Oceanography in Qingdao for its deep-sea submersible, Jiaolong.

"It will help, for instance, the simulation of the oceans on our planet with unprecedented resolution," Mr. Feng Liqiang, operational director of the Marine Science Data Center in Qingdao, told the *South China Morning Post*. "It will give China a bigger say over international affairs," he added, hinting at how the project aligns with the country's ambitions of becoming a maritime superpower.

JAPAN TO UNVEIL AI-CENTRIC SUPERCOMPUTER IN APRIL 2018

The AI Bridging Cloud Infrastructure (ABCI) at Japan's National Institute of Advanced Industrial Science and Technology is scheduled to begin operations in the second quarter of 2018. Set to be Japan's fastest supercomputer and potentially the third fastest in the world, the ABCI is intended to support artificial intelligence (AI) projects by governments and businesses.

The system will offer a theoretical peak performance of 37 petaFLOPS of double-precision performance, or, more importantly for AI users, 550 petaFLOPS of half-precision performance. Unlike scientific simulations which typically require double-precision or 64-bit numbers, deep learning applications use single-precision (32-bit) or half-precision (16-bit) numbers.

In line with its focus on AI applications, the ABCI will come equipped with a total of 4,352 Tesla V100 graphics processing units (GPUs), NVIDIA's machine learning accelerators.

"To facilitate high-performance scalable machine learning, ABCI utilizes supercomputer technologies to accommodate bandwidth in memory, interconnect and I/O to match the immense FLOPS it has, a factor that distinguishes it from normal cloud infrastructures," said Professor Satoshi Matsuoka, who leads the ABCI project hosted at the National Institute of Advanced Industrial Science and Techology AI Research Center (AIST-AIRC).

AURORA TO BE FIRST US EXASCALE COMPUTER

Originally scheduled to be delivered in 2018 as a 180-petaFLOP system, the Aurora supercomputer at the Argonne National Laboratory in the US will instead be ready by 2021 as a more powerful exascale system.

If completed according to schedule, Aurora will bring the country's exascale goal forward by two years, ahead of the 2023 target first announced in the Department of Energy's Exascale Computing Plan (ECP) in 2016. This, however, is still behind roadmaps issued by Japan, China and France, which all plan to have their exascale systems online by 2020.

Although the ECP largely escaped heavy cuts to the 2018 science and technology budget, the Aurora supercomputer was left in limbo when it was omitted from the Department of Energy's budget request. The current expansion of the scale of the project puts Aurora in the lead to be the US's first exascale system.

First announced in 2014 as part of the Collaboration of Oak Ridge, Argonne and Livermore (CORAL) initiative, Aurora was intended to be one of three 'pre-exascale' computers. Unlike the other two systems under CORAL—Summit at Oak Ridge National Laboratory and Sierra at the Lawrence Livermore National Laboratory—which are based on IBM and NVIDIA hardware, Aurora has been contracted to Intel and Cray.

Digital Dispatch

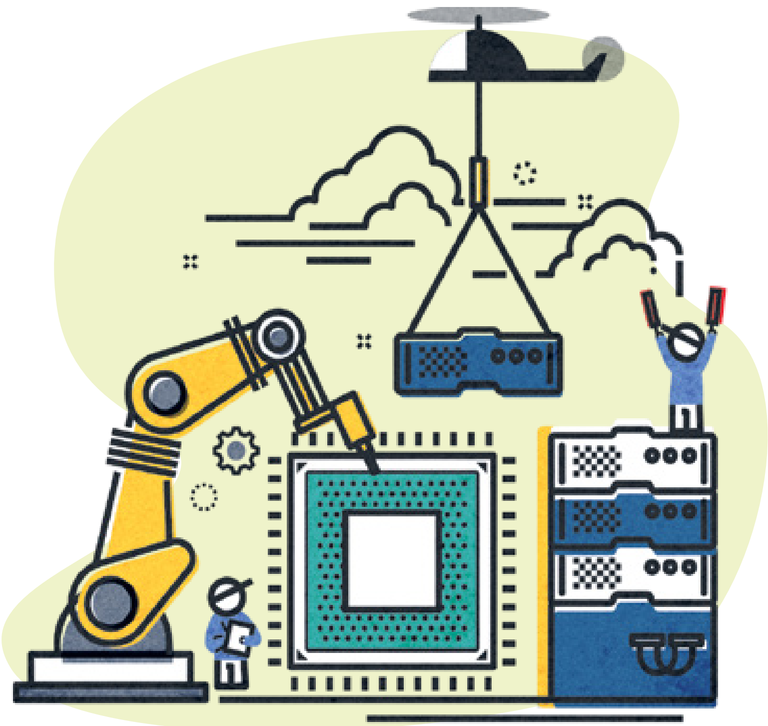

FIRST OF INDIA'S 70 SUPERCOMPUTERS ANNOUNCED

The Indian Institute of Technology, Kharagpur (IIT-KGP) is slated to become the first institution in India to run a supercomputer under the country's National Supercomputing Mission (NSM).

First announced in 2014 and approved in 2015, the NSM is a seven-year plan supported by Rs 4,500 crore (~US$700 million) in government funding.

"The new system, with both CPU and CPU-GPU based servers, along with the already existing HPC equipment, will provide about 1.5 petaFLOPS capacity support to several areas where the researchers of IIT-KGP are actively involved," said IIT-KGP's director, Prof P. P. Chakrabarti.

India currently has four supercomputers on the TOP500 list, the most powerful of which is the 901-teraFLOP system installed at the Supercomputing Education and Research Center at the Indian Institute of Science in Bangalore. Intended to jumpstart India's indigenous supercomputing capabilities, the NSM was established with the aim of helping India "leapfrog to the league of world-class computing power nations," according to a 2015 press release by the Government of India.

The three-phase strategy of the NSM involves first importing six supercomputers, followed by developing major parts such as switches and nodes, and finally manufacturing an entirely 'Made in India' supercomputer.

SINGAPOREAN STUDENTS SCORE RECORD-BREAKING WIN AT SC17

A team from Nanyang Technological University (NTU), sponsored by Singapore's National Supercomputing Centre, has walked away as champions of the Student Cluster Competition at SC17, a 48-hour challenge that gives teams the opportunity to design and build their own cluster computing systems.

Not only must teams build and optimize a supercomputer, they must also demonstrate its performance by running both the LINPACK and HPCG benchmarks, as well as a series of applications with real-world datasets. In addition, the SC17 challenge involves one or more power shutdowns that could occur anytime within the 48-hour live event, mimicking the scenario of the first student cluster challenge in 2007 when the power at the convention center suddenly failed.

The NTU team beat 15 other finalists, including crowd favorite Tsinghua University, which was hoping to score a 'triple crown' victory by winning at SC17. The Tsinghua team had won the Asia Student Supercomputer Challenge and ISC 2017, and were denied the opportunity of scoring a hat trick.

The winning team's risky decision to briefly switch off their fans during the LINPACK run paid off, giving them a record-breaking performance of 51.77 teraFLOPS per second. Impressively, the team also had the best HPCG performance, almost doubling the previous record.

WHAT'S UP!

SINGAPORE TO HOST HPC LUMINARIES AT SUPERCOMPUTINGASIA 2018

Building on the momentum generated over three successful runs of its flagship Supercomputing Frontiers conference, the National Supercomputing Centre of Singapore is hosting SupercomputingAsia 2018 (SCA18), an expanded conference that will feature five co-located allied events in addition to workshops, exhibitions and scientific presentations.

Keynote speakers of the inaugural SCA18 include Dr. Goh Eng Lim, vice president and chief technology officer of Hewlett Packard Enterprise; Dr. Satoshi Sekiguchi, vice president and director general of the Department of Information Technology and Human Factors at the National Institute of Advanced Industrial Science and Technology; Professor Joe Mambretti, director of the International Center for Advanced Internet Research, Northwestern University; and Professor Thomas A. DeFanti, a research scientist at Qualcomm Institute and distinguished professor emeritus at the University of Illinois.

Participants can also look forward to interacting with academic and industry thought leaders at co-located events such as the Asia-Pacific Advanced Network Meeting (APAN45) and the Conference on Next Generation Arithmetic (CoNGA). Awards will also be given out at SCA18 in recognition of the most outstanding HPC projects in science, industry and innovation across Asia.

For more information, visit sc-asia.org

WHAT

SupercomputingAsia 2018

WHEN

March 26–29, 2018

WHERE

Resorts World Convention Centre, Singapore

ISC HIGH PERFORMANCE RETURNS TO FRANKFURT FOR 2018

The ISC High Performance 2018 (ISC 2018) conference will be held in Frankfurt, Germany, from June 24-28, 2018. Previously known as the International Supercomputing Conference and widely considered to be the European counterpart of the SC Conference series held in North America, ISC 2018 is where the June edition of the TOP500 list is traditionally announced.

This year's conference will focus on 13 key topics, including exascale systems, future applications for quantum computers and artificial intelligence on HPC platforms. The topics were selected by a program committee chaired by Professor Horst Simon, two-time Gordon Bell Prize winner and deputy laboratory director and chief research officer of Lawrence Berkeley National Laboratory.

Highlights include the highly anticipated student cluster competition and the closely watched HPC in Asia country updates, chaired by Professor Lu Yutong, director of the System Software Laboratory at the National University of Defense Technology in China.

ISC 2018 will also feature workshops, tutorials, talks and poster sessions, as well as birds-of-a-feather sessions where like-minded participants can come together to discuss ideas and network.

For more information, visit isc-hpc.com

WHAT

ISC High Performance 2018

WHEN

June 24–28, 2018

WHERE

Frankfurt, Germany

RIDING THE HPC WAVE

Supercomputers in the search for oil

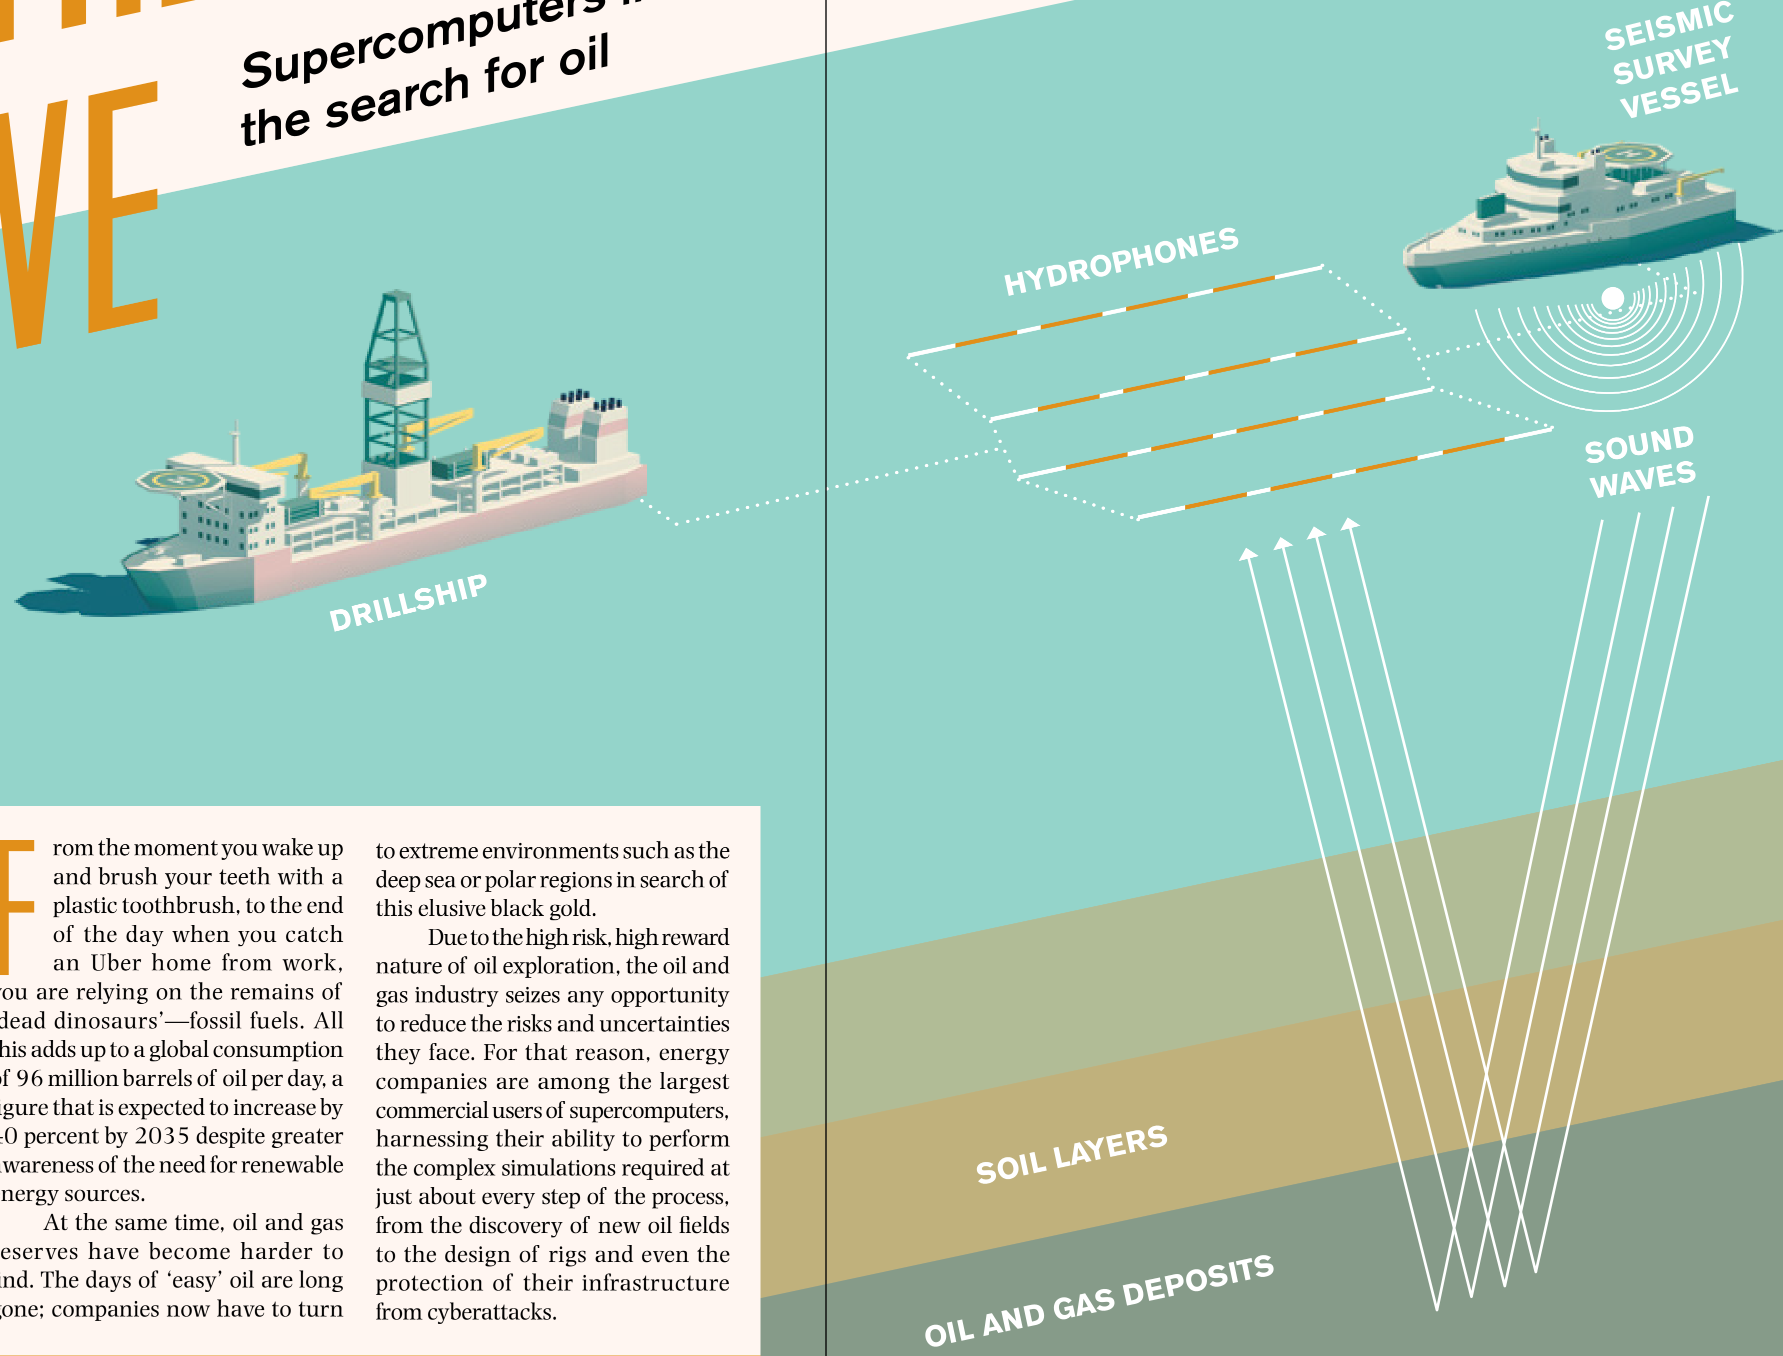

Supercomputers are useful in practically every stage of the oil production process. Here are five ways they play a role in the oil and gas industry.

By **Rebecca Tan**

From the moment you wake up and brush your teeth with a plastic toothbrush, to the end of the day when you catch an Uber home from work, you are relying on the remains of 'dead dinosaurs'—fossil fuels. All this adds up to a global consumption of 96 million barrels of oil per day, a figure that is expected to increase by 40 percent by 2035 despite greater awareness of the need for renewable energy sources.

At the same time, oil and gas reserves have become harder to find. The days of 'easy' oil are long gone; companies now have to turn to extreme environments such as the deep sea or polar regions in search of this elusive black gold.

Due to the high risk, high reward nature of oil exploration, the oil and gas industry seizes any opportunity to reduce the risks and uncertainties they face. For that reason, energy companies are among the largest commercial users of supercomputers, harnessing their ability to perform the complex simulations required at just about every step of the process, from the discovery of new oil fields to the design of rigs and even the protection of their infrastructure from cyberattacks.

1 EXPLORATION

More than 70 percent of the Earth's surface is covered in water, so quite naturally, oil companies need to search the seas as well as the land. Peering past the watery depths, however, is an extremely complicated affair. To visualize what lies beneath the ocean floor, companies use a technique called seismic imaging, sending sound waves into the ground using a device called an air gun, and collecting the sound waves that bounce back using a device called a geophone.

Each geophone collects data from several thousand sampling points, while each air gun used to generate the sound waves is in turn linked to thousands of geophones. Supercomputers are absolutely essential to help companies make sense of this large amount of raw data—which easily reaches hundreds of terabytes—and turn it into an actionable competitive advantage.

Seven or eight of the largest publicly owned oil and gas companies are known as supermajors, and collectively known as Big Oil. The Pangea supercomputer owned by French supermajor Total, for example, is used by the company to improve the accuracy of their subsurface imaging and save time by reducing the need for further exploration. Pangea is the world's largest industrial-use supercomputer, and with a peak performance of 6.7 petaFLOPS comes in at a respectable 21st position on the TOP500 list.

2

PRODUCTION

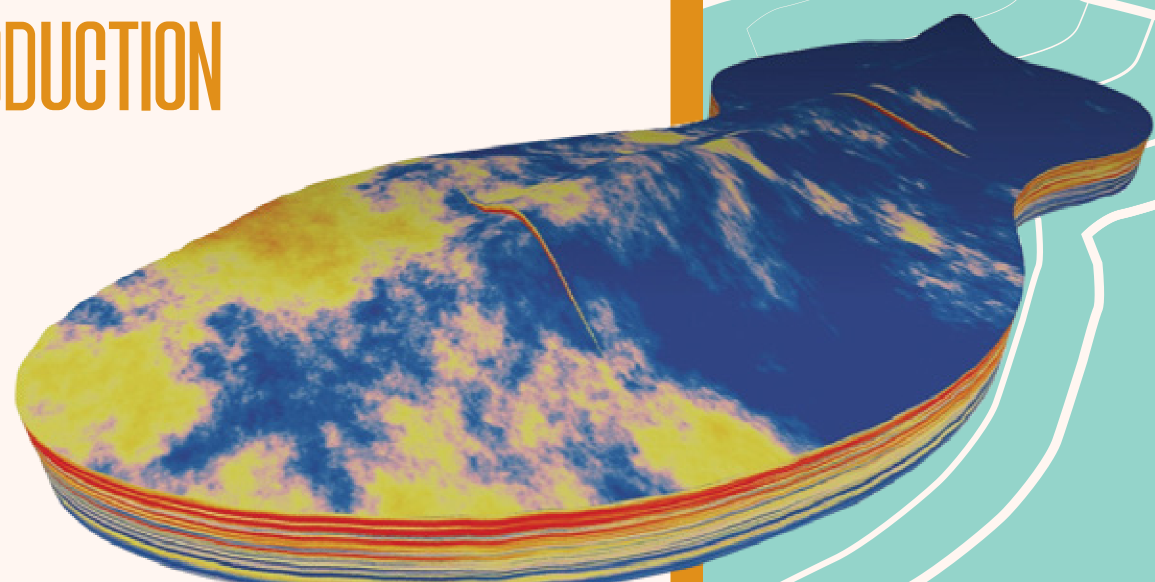

Visualization of a billion-cell petroleum reservoir model. Such highly detailed reservoir models make oil drilling and production more efficient. Credit: IBM.

After seismic imaging helps to locate a potential oil reservoir, the reservoir is extensively mapped using computer simulations designed to help engineers make decisions such as where to place oil wells and how to design facilities.

Reservoir modeling must take into account the complex flows of oil, water and natural gas, as well as the porosity and permeability of the surrounding rock. In the simulation, the entire reservoir is represented as a series of cells, with each cell having a set of values for measurements such as pressure and the relative concentrations of oil, water and gas. The more cells a reservoir model has, the higher its resolution and accuracy.

In February 2017, US oil giant ExxonMobil announced that they had successfully run a billion-cell reservoir model on the Blue Waters supercomputer at the National Center for Supercomputing Applications, using over 700,000 processor cores.

Just two months later, US technology company IBM announced that they had also run a billion-cell simulation, but this time using only 60 processors and 120 graphics processing units (GPUs), reducing the simulation run time from 20 hours to just one-and-a-half hours.

3

PROCESSING

Supercomputers don't stop being useful once the oil is found; they remain a vital part of daily operations even after the oil has been brought to the surface. One way that supercomputers are used in processing is in the design of separators, a piece of equipment that is used to separate the oil, water and gas into different streams for downstream processing.

In 2017, the Saudi Arabian Oil Company (Aramco) worked with engineering simulation specialist ANSYS to simulate a multi-phase gravity separation vessel. Although more realistic than individual computational fluid dynamics analyses, multi-phase models like the one used by ANSYS are more complex and difficult to scale.

With the help of the Shaheen II supercomputer at the King Abdullah University of Science and Technology, the team nevertheless managed to reduce the time taken for the simulation from several weeks to an overnight run.

Aramco, the largest oil and gas company on the world, said that this kind of simulation helps them reduce development time and allows operators to better predict the performance of equipment under varying conditions. Poor separation can reduce the productivity of an oil well by as much as 50 percent, making separator design critical for improving yields.

4
DESIGNING RIGS

Offshore rigs are significantly more expensive than land-based rigs; a basic offshore rig costs about US$200 million while rigs designed for more challenging environments can go up to almost a billion dollars each. Aside from being eye-wateringly expensive, offshore rigs are an engineering marvel, boasting thousands of kilograms of steel that somehow remain in the right position while floating in the middle of the ocean.

Floating rigs need to withstand extreme weather conditions and remain stable even in the face of strong ocean currents. In particular, vortex-induced motion poses a safety risk, adversely impacting the mooring systems of the offshore drilling system. As vortex-induced motion is difficult to reproduce in the lab, it must be studied through simulation instead.

In 2015, a team of researchers from the Los Alamos National Laboratory in the US used supercomputers to perform a comprehensive computational fluid dynamics analysis of vortex-induced motion, testing different turbulence models and confirming their model with experimental data. Their findings could be used to improve rig safety, as well as inform the design of other large floating structures.

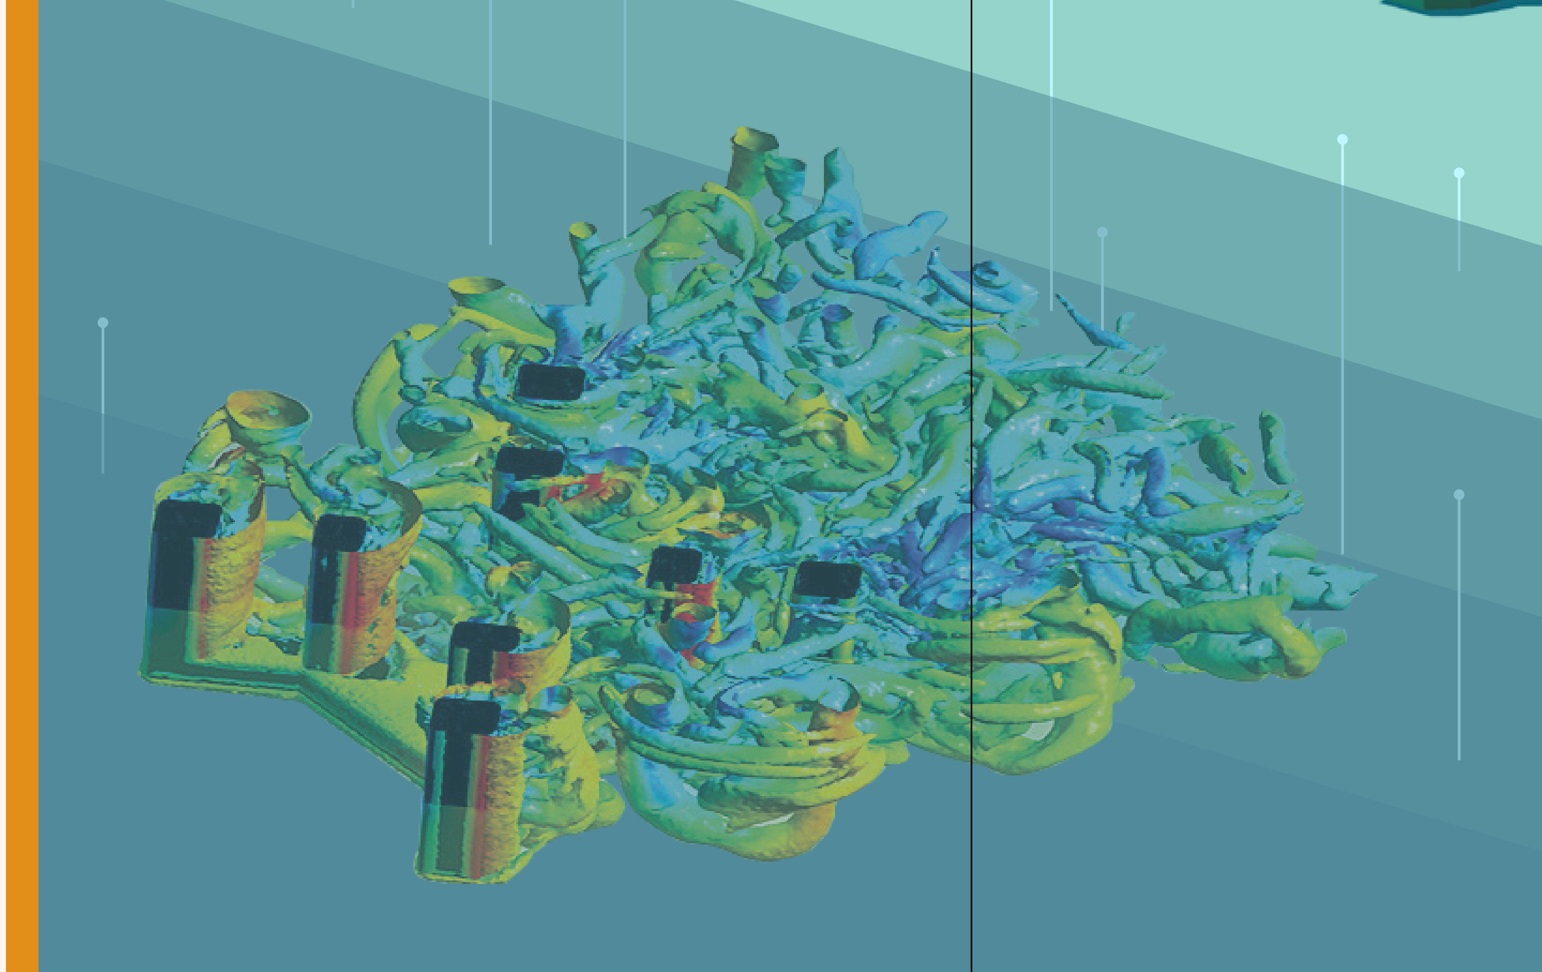

A simulation of vortex-induced motion showing how ocean currents affect offshore oil rigs. Credit: Los Alamos National Laboratory.

5
CYBERSECURITY

For as long as valuable goods have been transported over the sea, piracy has been a problem. According to a report by UK-based think tank Chatham House, Nigeria loses about US$1.5 billion a month to piracy and theft targeting the oil industry. These days, however, oil companies have to contend with much more than physical attacks by machine gun-toting pirates; they also have to worry about cybersecurity.

Three quarters of energy companies surveyed by the US-based privacy and data protection think tank Ponemon Institute experienced at least one cyberattack in 2016, and the energy industry is second only to the financial industry in terms of being prone to cyberattacks.

Alongside more traditional security measures, supercomputers can also play a role. For example, US-headquartered HPC provider Cray offers an analytics platform that combines supercomputing with enterprise-standard security frameworks to accelerate the detection of threats in cyberspace.

SUPERCOMPUTERS TAKE TO THE SEA

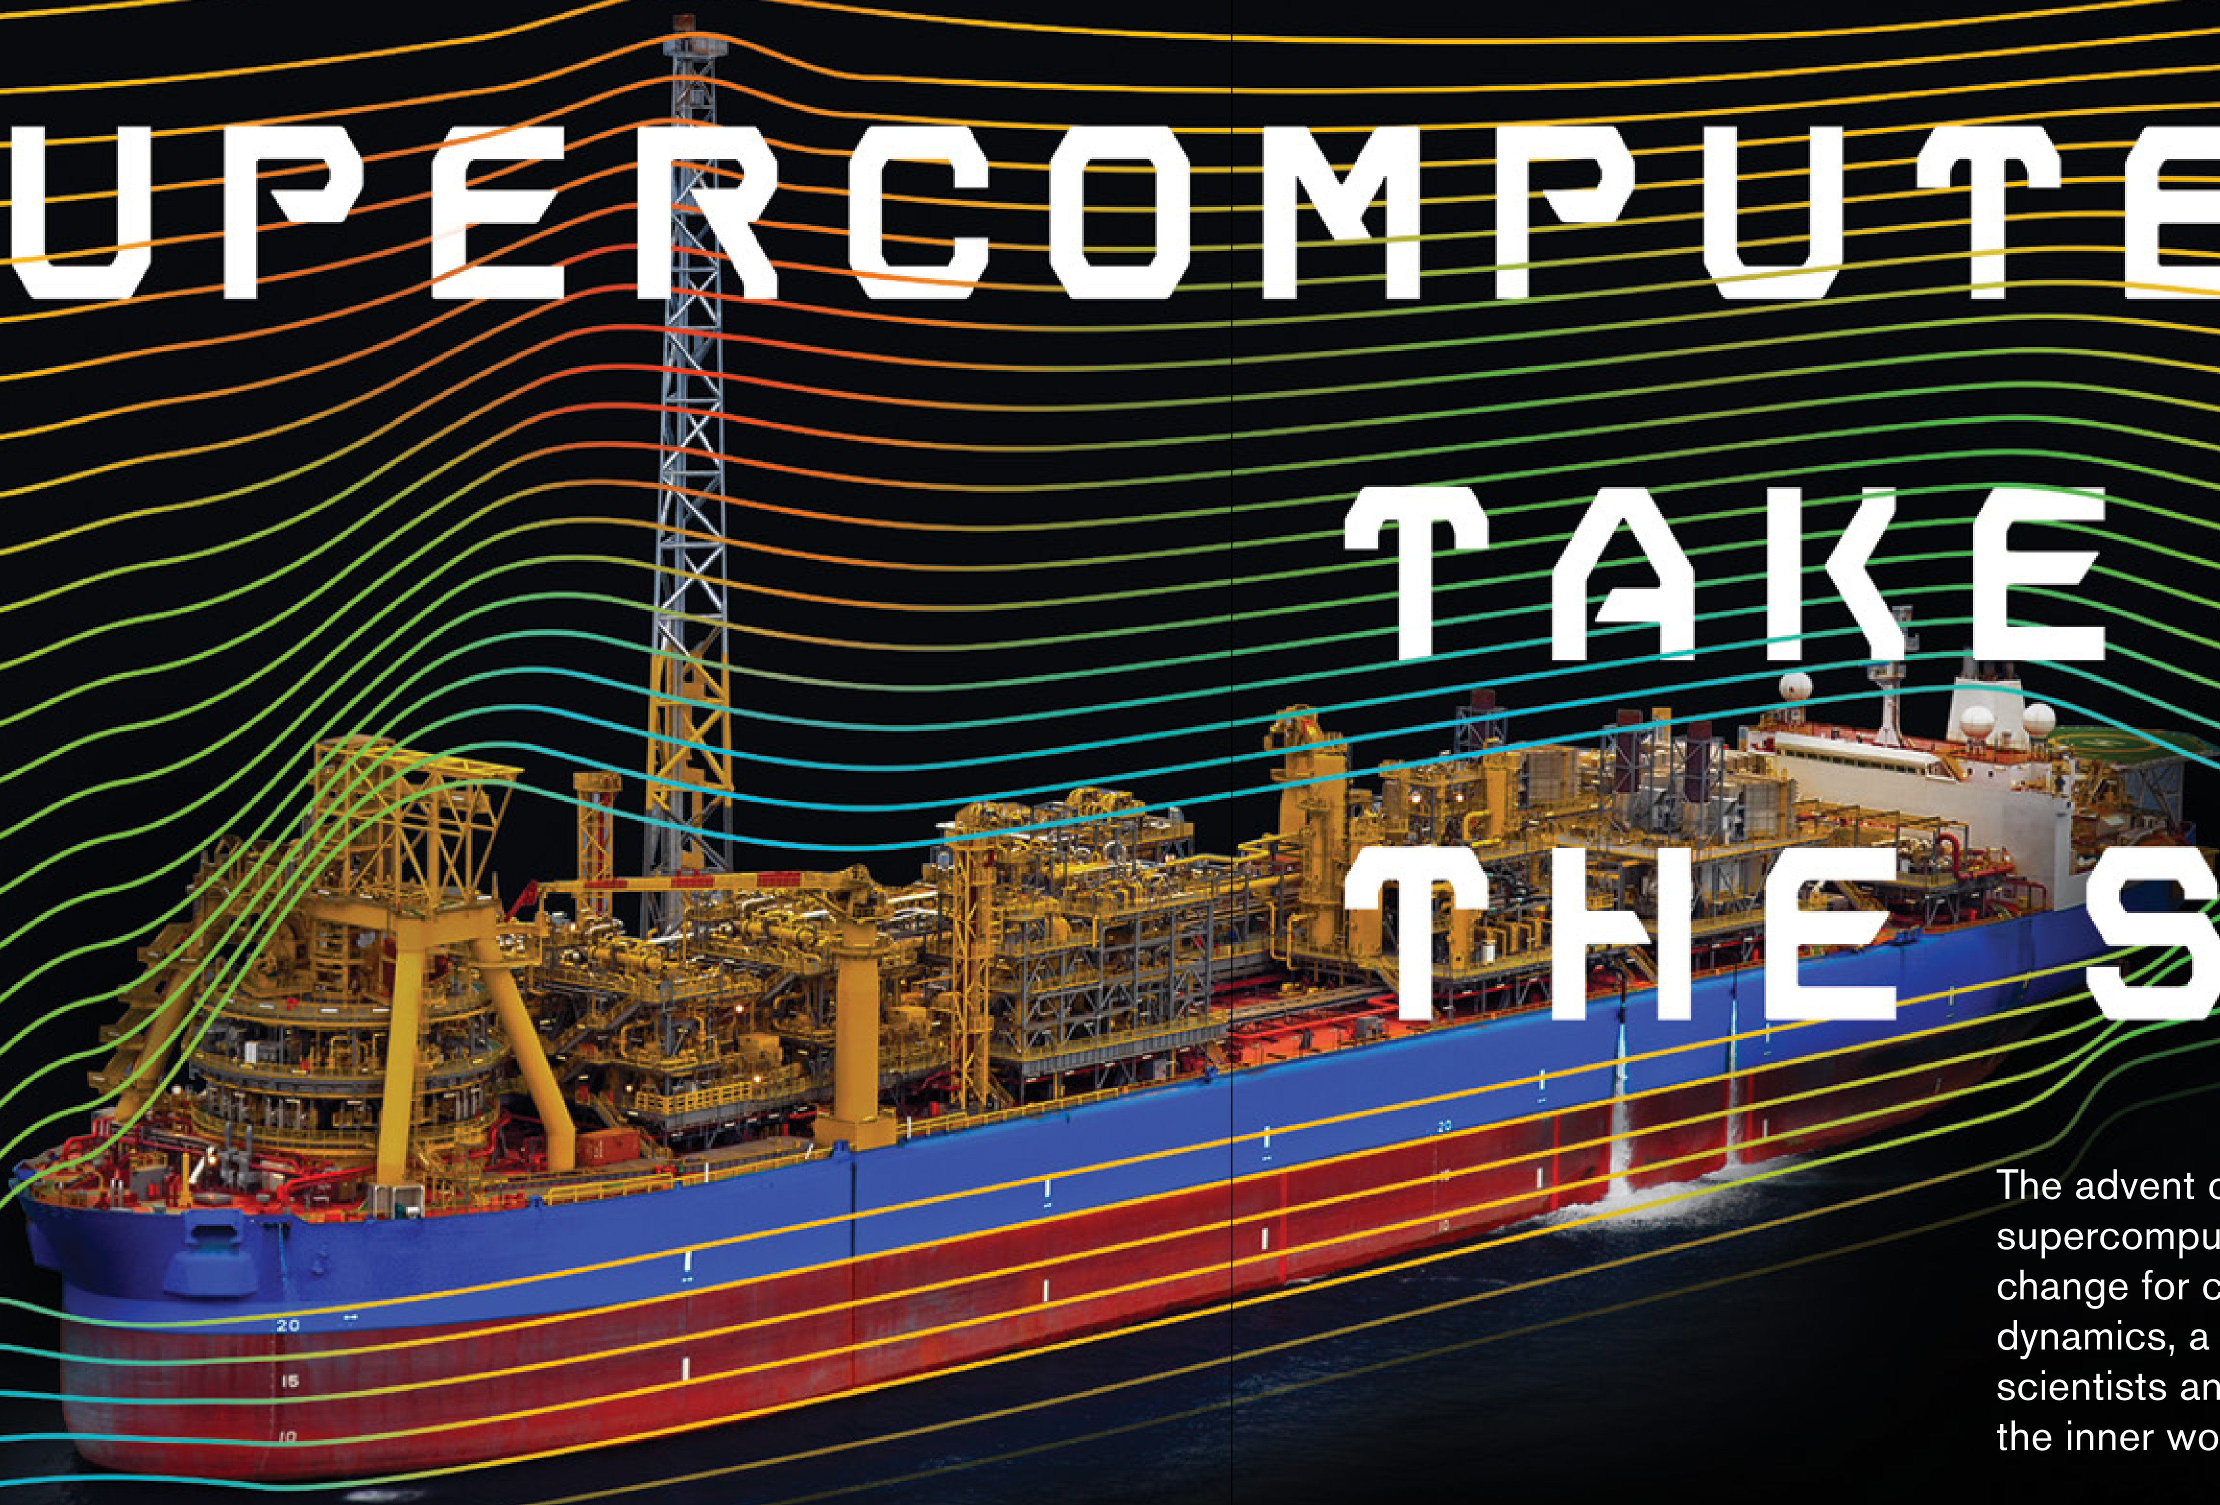

The advent of powerful supercomputers represents a sea change for computational fluid dynamics, a field which is helping scientists and engineers demystify the inner workings of the ocean.

By **Sim Shuzhen**

"THERE IS, ONE KNOWS NOT WHAT SWEET MYSTERY ABOUT THIS SEA, WHOSE GENTLY AWFUL STIRRINGS SEEM TO SPEAK OF SOME HIDDEN SOUL BENEATH . . . ."

Herman Melville, *Moby-Dick*

FOR CENTURIES, MAN HAS BUILT ALL MANNER of ships, submersibles and structures to explore the breadth and depth of the ocean and harvest its bounty. Yet—as captured in Herman Melville's description of the Pacific Ocean in the novel Moby-Dick—the wild blue yonder continues to both fascinate and confound us with its vast, mercurial nature:

"There is, one knows not what sweet mystery about this sea, whose gently awful stirrings seem to speak of some hidden soul beneath . . . ."

Whalers and sailors may be physically buffeted about by the ocean, but physics, too, has a hard time dealing with its unpredictability. The field of fluid dynamics, which describes the flow of fluids—in this case, the motion of water within the ocean—is governed by mathematical equations that are notoriously difficult to solve. Researchers have thus resorted to numerical simulations to predict the behavior of currents, waves and eddies.

Complex behavior begets complex simulations; these are so computationally intensive that it is only recently, with the advent of powerful supercomputers, that researchers have been able to model fluid dynamics faster and more accurately.

With their newfound ability to read the seas, supercomputers are now being put to work across a wide range of marine-related research efforts—modeling global ocean climate change, understanding the hydrodynamics of coral reef ecosystems, and even sculpting perfect waves for professional surfers, just to name a few.

WEATHERING A DOWNTURN

One group with enormous commercial interests in these recent advances is the offshore and marine industry—a broad sector that includes oil and gas companies, shipyards, and firms that engineer and build oil rigs, platforms and other offshore structures. The industry has been in the doldrums since oil prices crashed in mid-2014; since then, nearly US$400 billion worth of proposed energy projects worldwide have been put on hold, estimates UK-based energy research and consulting firm Wood Mackenzie.

"At the moment it is still not very clear when oil prices are going to pick up significantly, so the [offshore and marine] market as a whole has slowed down in many ways," said Professor Chan Eng Soon, CEO of the Technology Centre for Offshore and Marine, Singapore (TCOMS), a research and development centre co-founded by the Agency for Science, Technology and Research and the National University of Singapore in 2016 to help the industry take advantage of emerging digital technologies.

"On the other hand, because of this, there is a push for technology and solutions to disrupt current practices where possible, and in more cost-effective ways," he said.

Singapore has in many ways always been a maritime nation, added Chan. The country has one of the busiest ports in the world, and commands about 70 percent of the global market for both jack-up rigs and floating production storage and offloading (FPSO) platforms. In 2014, annual output from its offshore and marine sector was nearly S$25 billion (~US$18.6 billion), making up nearly two percent of the country's gross domestic product.

But depressed oil prices have since taken their toll: output from the sector in October 2017 was less than half what it was in the same month three years ago, and Singapore companies Keppel Corp and Sembcorp Marine, the world's largest builders of oil rigs, have seen share prices and earnings fall.

A DIGITAL REVIVAL

The future looks less bleak—with oil prices inching up towards US$60 a barrel as of December 2017, the industry is showing signs of rallying. Still, weathering the downturn has made companies more eager than ever to make the most of digital technologies. Supercomputing is a large part of this push, particularly for projects involving computational fluid dynamics, such as rig and vessel design.

"Conventional design approaches require significant resources and testing of physical models, and are thus time-consuming and expensive. Critically, knowledge and skills are developed in isolation and are non-transferable," said Mr. Aziz Merchant, executive director of the Keppel Offshore & Marine Technology Centre (KOMtech), Keppel O&M's research and development arm, in an interview with *Supercomputing Asia*.

Keppel O&M is drawing on the National Supercomputing Centre Singapore's (NSCC) one-petaFLOPS ASPIRE 1 supercomputer to run advanced computational fluid dynamics simulations, which numerically capture the environmental loads on vessels and rigs, as well as how these structures respond.

This data has allowed the company's engineers to optimize their designs and develop innovative new ones, said Merchant. Keppel O&M has, for example, developed technologies that reduce motion and improve safety on semisubmersibles and accommodation rigs; unique 'ice-class' vessels that can access frozen areas of the ocean; and optimized, low-resistance hull forms that make supply vessels more fuel-efficient.

SIMULATING THE SEA

TCOMS is also in the business of enhancing simulation capabilities. A key feature of the Centre is its deepwater ocean basin facility, slated for completion in 2019. Using wave generators and supercomputer-powered computational fluid dynamics, the facility's 50-meter-deep central pit can simulate marine environments down to 3,000 meters, said Chan. For comparison, the world's deepest oil and gas project, Shell's Stones, operates in some 2,900 meters of water in the Gulf of Mexico.

Despite its name, the basin can also be used to model shallower waters, added Chan. Researchers can thus study the hydrodynamics that surface vessels—those that lay pipes or supply the rigs and platforms, for example—are subject to, and use the information to develop improved or autonomous versions.

But in order to build truly detailed simulations—digital twins—of the rigs or vessels being studied, other types of input are also needed. In addition to parameters from the basin, real-world data on waves and currents, collected by sensors mounted on vessels and rigs, can also be weaved into the simulacra.

"The oil and gas industry has been pushing the use of big data for a number of years, with many groups pursuing sensing, data analytics and deep learning, all geared towards the overall concept of digitalization," said Chan, adding that TCOMS is keen to partner with industry players to gather data and develop solutions with real-world applications.

Another goal is to be able to understand the behavior of more complex, non-linear systems—an ecosystem of rig, platform, vessels and submersibles, for example.

"What has evolved out of digitalization is the ability to address complexities at the system-of-systems level," Chan explained. "The behavior of complex systems tends to be very non-linear and to some extent still unpredictable, so there are technical challenges in that sense. We have embarked on the journey to push boundaries and advance the state-of-the-art."

When it comes to supercomputing, the marine and offshore industry in Singapore is a late bloomer as compared to early adopters in finance, aerospace, healthcare and other sectors, said Merchant. "The marine and offshore industry recognizes that it is lagging behind adjacent industries and needs to ramp up quickly," he noted.

Chan is optimistic that supercomputing initiatives will help move the industry forward. "We are focused on technologies that are geared towards future smart systems. We are not doing more of the same; instead, we are looking at how we can deepen our fundamental knowledge and derive innovative solutions in partnership with the industry," he said.

INSTEAD OF TITANIC-SIZED OCEAN LINERS, THE VOLATILITY OF THE OCEAN, IT SEEMS, MAY BE BETTER TAMED BY LINES OF CODE.

THE STATE OF THE OCEAN

Supercomputers can do more than help mine the ocean for resources; they also help oceanographers and climatologists understand its workings on a grander scale. One of these researchers is Dr. Shuhei Masuda, a group leader and senior scientist who studies ocean circulation at the Japan Agency for Marine-Earth Science and Technology's (JAMSTEC) Research and Development Center for Global Change.

By storing solar radiation and helping to distribute heat around the globe, the ocean plays a crucial role in keeping the planet warm. To better understand these global phenomena, Masuda builds computational models of the ocean state—the overall circulation and climate patterns of huge bodies of water.

"Our main goal is to comprehend past and current ocean states, and to clarify the mechanism of ocean climate changes. This leads to a better understanding of the systems which gently regulate the Earth's climate," said Masuda. "Such knowledge is definitely required when making a future projection on the state of the global climate, for instance, in conjunction with global warming."

Masuda draws on a wide variety of ocean observations made by ships, moored buoys and floats; one of his resources is the Argo array of nearly 4,000 floats which gathers data on temperature, salinity and ocean velocity. Argo is run by an international consortium of climate research institutions, including JAMSTEC.

Photo credit: JAMSTEC

Earth Simulator

JAMSTEC's Earth Simulator supercomputer is used for various marine-earth research applications, including studying ocean circulation and the evolution of the Earth's interior, as well as making global warming projections.

Photo credit: JAMSTEC

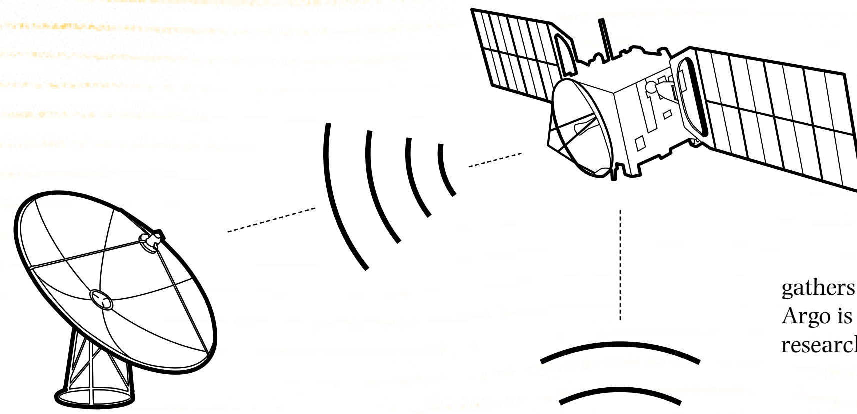

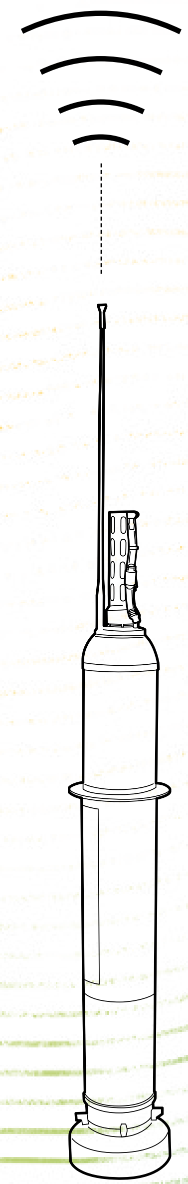

MIRAI

The oceanographic research vessel helps JAMSTEC researchers observe water temperature, salinity and tidal current, and collect sediments on the seabed.

Argo floats

Argo floats are autonomous drifting buoys that move up and down in the ocean from the sea surface to 2,000-meter depths, measuring water temperature and salinity.

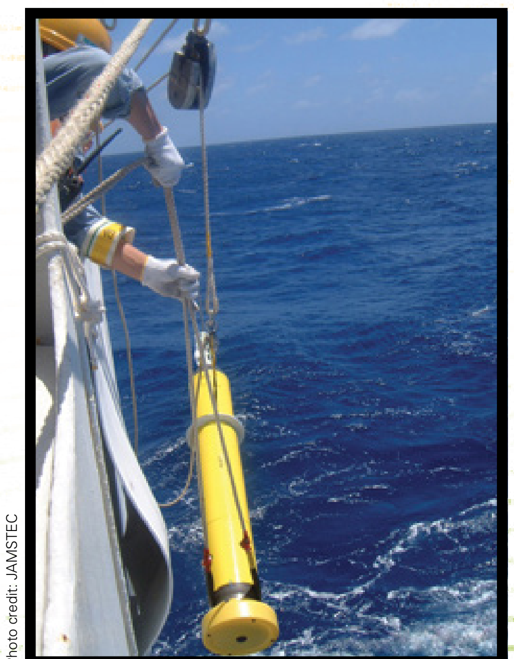

Photo credit: JAMSTEC

A SEA OF CODE

On the aptly named Earth Simulator, JAMSTEC's terascale supercomputer, Masuda synthesizes real-world data points into computational models to make estimates of the ocean state. These can then be used to study how ocean climate has changed over time, and to make projections of what it will be like in the future, he explained.

Masuda has used computational models to suggest why a major, widely predicted El Niño event failed to materialize in 2014. Existing models did not take into account longer-term, five-to-ten-year variations in tropical seasonality, which affect the amount of heat energy transferred to the oceans, he and his collaborators found. By incorporating these variations into ocean state estimates, the researchers were able to accurately 'hindcast' past El Niño events, and also—hopefully—better predict future ones.

JAMSTEC now houses a new supercomputer at its Yokohama Research Institute location. At 19 petaFLOPS and 20 million cores, the Gyoukou supercomputer—co-developed by ExaScaler and PEZY Computing, and unveiled in November 2017—now sits in fourth position on the TOP500 ranking of the world's fastest supercomputers.

In Asia, Singapore and Japan are far from alone in their efforts at using supercomputers to describe and predict the behavior of the ocean. Twin petascale supercomputers Miri and Nuri power the Korea Meteorological Administration's climate research; meanwhile, in China, an exascale supercomputer dedicated to analyzing ocean data and expanding the country's maritime presence may be ready as early as 2019.

These new levels of computing power will enable researchers to develop higher-resolution models that can project patterns out over longer time frames, said JAMSTEC's Masuda. Instead of Titanic-sized ocean liners, the volatility of the ocean, it seems, may be better tamed by lines of code.

THE VIEW FROM THE TOP500

Tracking supercomputers with LINPACK and beyond

The TOP500 list may have begun almost as an accident, but it has nonetheless come to shape the way we think about supercomputers. *Supercomputing Asia* catches up with the co-originator of the list, Jack Dongarra, to find out more about its past, present and future.

By **Rebecca Tan**

In the relatively staid world of supercomputing, there is nothing quite like the biannual announcement of the TOP500 list, the closest thing the community has to its own version of the Olympics. In June and November every year, the ranking of the world's top supercomputers is released to an audience of thousands, including scores of journalists who instantly begin tweeting the results to those unlucky enough to miss out on the event.

For the last 25 years, the closely watched TOP500 list has kept the world informed about the latest developments in the fast-moving field of high performance computing (HPC). It might come as a surprise, then, to find out that for Professor Jack Dongarra benchmarking is but a hobby.

"The whole benchmarking thing came about almost as an accident," Dongarra told *Supercomputing Asia*. "When I'm talking to somebody outside of my field and they ask what hobbies I have, I would say, 'I benchmark supercomputers'."

UNPACKING THE LINPACK STORY

Computers on the TOP500 list are ranked according to how quickly they can solve a set of problems in the LINPACK software library, which, as the name suggests, is a package of linear algebra problems. LINPACK was developed over the course of three intense summers, where Dongarra, Cleve Moler, Pete Stewart and James Bunch met at the Argonne National Laboratory to thrash out what would eventually be included. In 1979, the group finally released LINPACK to the world, along with a manual to show people how to use the software.

Tucked into the back of that manual, in what first began as a series of hand-written scribbles, was an appendix that became the very first LINPACK benchmark.

"It was a table of performance numbers for one of the routines in LINPACK that I had the opportunity to run on maybe close to 20 different computers," said Dongarra, who is currently the director of the Innovative Computing Laboratory at the University of Tennessee in the US. "The table grew as I maintained it over the years, adding to it whenever somebody came to Argonne and wanted to sell us a computer."

By the 1990s, Dongarra's list had hundreds of entries, allowing him to rank the different computers based on how they had performed on the LINPACK benchmark. Hans Meuer, a computer scientist at the University of Mannheim, Germany, maintained a similar list, ranking computers by their peak performance. In 1992, Meuer approached Dongarra suggesting that they should combine their lists, and thus the first TOP500 list was born.

FROM BENCHMARK TO BOOKENDS

In the quarter century since then, supercomputers have vastly improved. The most powerful supercomputer on the first TOP500 list, the Fujitsu Numerical Wind Tunnel at Japan's National Aerospace Laboratory, topped the list with a peak performance of 124 gigaFLOPS. In comparison, the most powerful computer on the November 2017 list was the Sunway TaihuLight, which at 93 petaFLOPS is almost a million times faster.

Despite these dramatic improvements, the LINPACK benchmark is still used to compare machines on the TOP500, and does not look set to be replaced any time soon. Some of the reasons the LINPACK benchmark has

Professor Jack Dongarra
University of Tennessee
Oak Ridge National Laboratory

has stood the test of time is that it is easy to run and understandable to a broad community, Dongarra said.

"But most importantly, it provides a historic reference point, giving us a good snapshot of supercomputing over the last 25 years. It allows us to look at trends and what the impact of various different architectures has been," he continued.

Nonetheless, LINPACK has its limitations. When it was designed nearly 40 years ago, floating-point arithmetic was very computationally expensive, making it the most important thing to optimize. "Today, our computers are over-provisioned for floating-point calculations; the more important thing now is data movement," Dongarra said.

To do numerical operations on data, the data first has to be moved through the memory hierarchy of the computer from the main memory, through to different levels of cache and then finally to the place where the arithmetic will be performed, the register. While modern processors can perform numerical operations very quickly—completing roughly 32 floating-point operations every cycle—it could take several hundred clock cycles to move the data from the memory to the register.

So instead of simply measuring how quickly computers can perform floating-point operations, a more accurate reflection of how a supercomputer handles real-world problems would be one that captures its ability to handle memory as well. The high performance conjugate gradient (HPCG) benchmark is one such measure, and since the November 2017 edition, has been featured on the TOP500 list alongside the LINPACK benchmark.

"We think of the LINPACK and HPCG as bookends, where LINPACK gives you a number that is very close to the peak performance of a system while the HPCG reflects the lower end," Dongarra explained. "Your own application—which is the best benchmark of performance—will fit somewhere between those two points, most likely towards the bottom."

The four individuals who developed the LINPACK software library: (L to R) Jack Dongarra, Cleve Moler, Pete Stewart and James Bunch in 1978, along with Dongarra's car and LINPACK license plate. Credit: Jack Dongarra.

TRACING THE TRENDS

Although the LINPACK benchmark no longer reflects how supercomputers are used today, Dongarra nonetheless feels that it provides a valuable perspective. "We don't want to lose that historic information, but we want to augment it with other kinds of benchmarks such as the HPCG benchmark."

By allowing us to compare different computer architectures over the decades, for example, the LINPACK benchmark reveals some interesting trends. In the early 90s, the list was dominated by vector-based machines which produced very high performance but were also very expensive. As cheaper, mass-produced or 'commodity' processors became more powerful, they gradually displaced vector computers and ushered in the era of parallel and distributed computing.

"Today, we see diverse architectures for HPC emerging. One is to use more and more commodity processors; the second is to offload the floating-point computation to an accelerator such as a graphics processing unit; and the third is to use what I call lightweight cores," Dongarra said.

As important as these different hardware architectures are, Dongarra stresses that the software ecosystem for scientific simulation and computational modeling must be developed in tandem with the investments in building machines.

"At the application level, the science has to be captured in mathematical models, which in turn are expressed algorithmically and ultimately encoded as software," he said. "This process also relies on a large infrastructure of mathematical libraries, protocols and system software that takes years to build up and must be maintained, ported and enhanced."

Tucked into the back of that manual, in what first began as a series of hand-written scribbles, was an appendix that became the very first LINPACK benchmark.

The scientific problems that supercomputers were designed to solve require close collaboration between domain-specific scientists, computer scientists and applied mathematicians, Dongarra continued. "To be able to run scientific applications on petascale systems, with tens of thousands of processors and extract all the performance that these platforms can deliver, demands that all parties involved work together to develop the necessary software."

SCALING SOFTWARE TO THE NEXT LEVEL

This brings us away from his hobby and back to Dongarra's 'day job,' where he is a principal investigator for three of the 35 software development projects funded by the US Department of Energy's Exascale Computing Project.

The first project, called Software for Linear Algebra Targeting at Exascale (SLATE), is a next-generation linear algebra program that Dongarra and his team are developing to run efficiently on exascale systems, supercomputers that are at least ten times faster than China's Sunway TaihuLight.

Parallel Runtime Scheduling and Execution Control (PaRSEC), on the other hand, is a program that will allocate tasks from SLATE to the hardware components available on a supercomputer, a job that is complicated because of the sheer number of ways to prioritize and execute the many tasks.

Last but not least is the Exascale Performance Application Programming Interface (EXA-PAPI), an interface that tracks diagnostic information such as the use of memory bandwidth to help users understand how well their software is performing on the hardware.

"These three projects tackle some of the most challenging—and interesting—problems standing in the way of exascale computing," Dongarra said.

"Software routinely outlasts—by years, and sometimes even decades—the hardware that it was originally designed to run on, as well as the individuals who designed and developed it. With so many problems to overcome, and the new ways of thinking that it has prompted, this is one of the most exciting times I have faced in my career," he concluded.

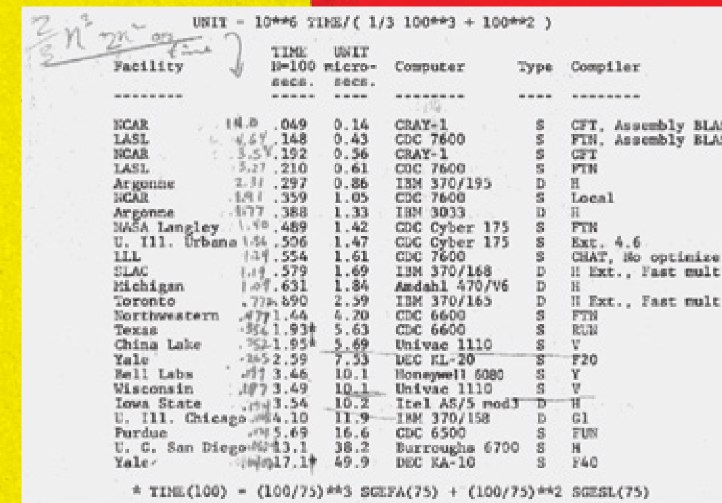

UNIT - 10**6 TIME/(1/3 100**3 + 100**2)

Facility	TIME N=100 secs.	UNIT micro-secs.	Computer	Type	Compiler
NCAR	.049	0.14	CRAY-1	S	CFT, Assembly BLAS
LASL	.148	0.43	CDC 7600	S	FTN, Assembly BLAS
NCAR	.192	0.56	CRAY-1	S	CFT
LASL	.210	0.61	CDC 7600	S	FTN
Argonne	.297	0.86	IBM 370/195	D	H
NCAR	.359	1.05	CDC 7600	S	Local
Argonne	.388	1.33	IBM 3033	D	H
NASA Langley	.489	1.42	CDC Cyber 175	S	FTN
U. Ill. Urbana	.506	1.47	CDC Cyber 175	S	Ext. 4.6
LLL	.554	1.61	CDC 7600	S	CHAT, No optimize
SLAC	.579	1.69	IBM 370/168	D	H Ext., Fast mult.
Michigan	.631	1.84	Amdahl 470/V6	D	H
Toronto	.690	2.59	IBM 370/165	D	H Ext., Fast mult.
Northwestern	1.44	4.20	CDC 6600	S	FTN
Texas	1.93*	5.63	CDC 6600	S	RUN
China Lake	1.95*	5.69	Univac 1110	S	V
Yale	2.59	7.53	DEC KL-20	S	F20
Bell Labs	3.46	10.1	Honeywell 6080	S	Y
Wisconsin	3.49	10.1	Univac 1110	S	V
Iowa State	3.54	10.2	Itel AS/5 mod3	D	H
U. Ill. Chicago	4.10	11.9	IBM 370/158	D	G1
Purdue	5.69	16.6	CDC 6500	S	FUN
U. C. San Diego	13.1	38.2	Burroughs 6700	S	H
Yale	17.1*	49.9	DEC KA-10	S	F40

* TIME(100) = (100/75)**3 SGEFA(75) + (100/75)**2 SGESL(75)

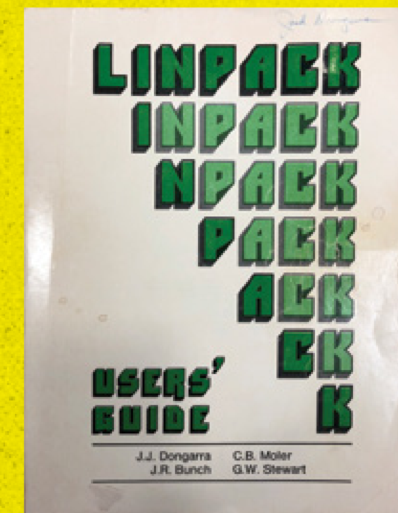

Dongarra's hand-written calculation of the megaFLOPS rate of various systems (top) and the cover of the first LINPACK user's guide in 1979 (right). Credit: Jack Dongarra.

COMPUTING IN THE COSMOS

A supercomputer that's out of this world

Supercomputing Asia talks to Hewlett Packard Enterprise's Dr. Goh Eng Lim about the company's supercomputer that is boldly going where no supercomputer has gone before.

By **Sim Shuzhen**

When dealing with the vast distances involved in interplanetary travel, Earth-based navigation systems can be off by hundreds of kilometers. But what if spacecraft could triangulate their own positions using signals from pulsars—neutron stars that emit radiation at highly regular intervals—just as sailors used to navigate by the heavens?

This is more than just conjecture—a pulsar-based navigation system is already being tested aboard the International Space Station (ISS). But this and other applications for deep space exploration are likely to run up against a major obstacle: the lack of compute power in space.

"If you travel to the moon, you can still rely on Earth-based computers to provide answers within seconds. But on Mars, it could take 20 minutes to transmit a message and another 20 minutes to receive an answer," Dr. Goh Eng Lim, vice president and chief technology officer, high performance computing (HPC) and artificial intelligence (AI), Hewlett Packard Enterprise (HPE), tells *Supercomputing Asia*. "As you travel further from Earth, you'll need to carry more computing resources with you."

Goh is the lead investigator on a joint HPE-NASA endeavor to build a supercomputer which can weather the harsh conditions of space. In September 2017, the Spaceborne Computer—a two-node machine weighing about 60 kilograms on Earth—was installed on the ISS, where it will undergo a year of rigorous testing.

THE APP STORE AT THE END OF THE UNIVERSE

In space, cosmic radiation wreaks havoc with computer circuitry, causing frequent glitches. Computers thus need to be physically 'hardened' or shielded against radiation. But by the end of this expensive, time-consuming process, they may be several generations behind the latest models.

Instead of hardware modifications, the Spaceborne Computer uses software to slow down its operations in order to prevent damage during a radiation event, meaning that it can be used straight out of the box.

Computers on space missions today also tend to have very specialized functions. "You have to think very carefully beforehand about what precious applications you intend to carry with you, because you will be embedding them in those hardened computers," says Goh. "If something new comes up, it's more difficult for them to run applications that are unanticipated."

But one-trick ponies are not ideal for the vagaries of a round trip to Mars. For longer, more complex space missions, having general purpose capabilities will minimize the need for astronauts to improvise in the event of an emergency.

The Spaceborne Computer, with its off-the-shelf hardware and Linux-based operating system, has the potential to be this all-rounder. Goh envisions space missions using it the way we use smartphones on Earth—for everything.

"Before launch, you can sit down and load all the applications you think you might need on this commercial off-the-shelf (COTS)-based computer. As you fly, if something unanticipated comes up, you hope you have preloaded the relevant app somewhere on the system," he explains. "This is a very powerful concept."

TRAVELING LIGHT

On the ISS, the Spaceborne Computer has achieved speeds of up to one teraFLOP—an order of magnitude more powerful than anything else in space, though still a far cry from its much larger Earth-bound counterparts.

But even if we could launch the world's fastest supercomputer into space, its power needs would quickly suck the spacecraft dry, says Goh. Traditional high-performance computing uses physics-based models to run simulations and make predictions—a top-down, computationally intensive approach.

Launching more compute power into orbit is thus not the only goal. What astronauts also need is a lightweight, versatile computing system that can help them make quick decisions in a pinch.

This is where machine learning comes into the picture, says Goh. Here, instead of traditional physics-based models, computer algorithms would learn by ingesting large amounts of data from past simulations, and then assign weights to a range of parameters according to their importance. These weights can then be used to make inferences, or predictions, about a given situation.

"This is what you carry with you when you travel—just the weights—and you can do lightweight predictions because you've done all the heavy lifting beforehand [on Earth]," Goh says. "I believe this is the approach we will quite often need for long duration space travel."

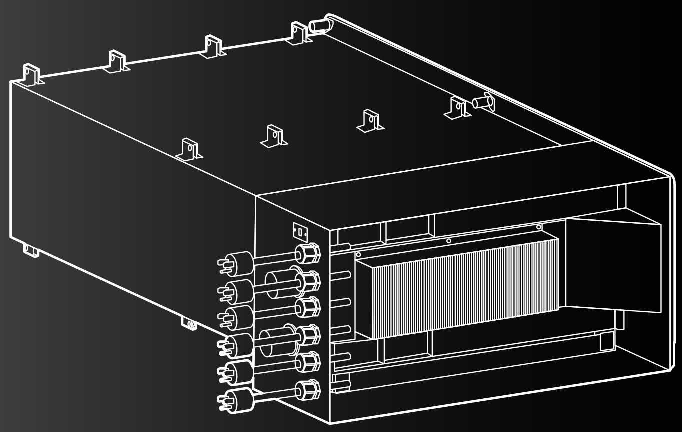

A rendering of the Spaceborne Computer, which is currently in tests at the International Space Station. Credit: NASA.

SCALING ASIAN SUPERCOMPUTING

Today, increasingly massive amounts of data are available for machine learning algorithms to ingest, meaning that they will have to do so more quickly. But many machine learning approaches, in particular deep neural networks, do not scale well, said Goh.

Back on Earth, the Tokyo Institute of Technology (TiTech) is working to address this very problem. "TiTech came to us with a need—to scale artificial intelligence (AI). It was a very clear, simple mandate," says Goh.

HPE built TiTech the Tsubame 3.0 supercomputer, a 47-petaFLOPS "massive bandwidth machine," as Goh calls it, with multiple tightly connected nodes that voraciously consume and crunch data. One of the fastest AI supercomputers in the world, its architecture also lays the groundwork for machine learning to be used in concert with more traditional high-performance computing approaches.

"Physics, of course, will need to be there as it always has, but you need to complement it with machine learning," says Goh. "Increasingly, we are seeing physics-based, top-down methods of prediction working side-by-side with bottom-up machine learning algorithms."

When it comes to pioneering new approaches in computer science such as this, Asia can rely on one advantage—its large number of STEM graduates, thinks Goh. "Smart people are everywhere, but to solve a problem you need a lot of them to come together. In Asia, there's still a very strong attraction for people to study STEM, and Confucian systems tend to give more reverence to education."

Still, whether Asian governments have the appetite to support bigger, more profound projects with no immediate economic returns remains an open question. Take the Laser Interferometer Gravitational-Wave Observatory (LIGO) in the US, for example—its Nobel prize-winning detection of gravitational waves was decades (and hundreds of millions of dollars) in the making.

"If the gravitational wave pitch were given worldwide, would any Asian government take that up? I think that's a good question to ask. The [US] National Science Foundation did, so credit to them—they also have their funding challenges, but it was profound enough for them to take it up," says Goh.

ADVANCES ON THE CARDS

In the future, supercomputers in space could benefit from an earthly pursuit—poker.

In Go and chess, players are privy to their opponent's position on the board. But poker is a different beast—not only do players not know their opponent's cards; they also go out of their way to mislead one another.

Incomplete information is often part and parcel of real-world situations—consider contract negotiations, auctions and military strategy, for example. The same is true for space travel. "On Earth, you can amass resources to get information to make decisions," says Goh. "But in space, you have less time and resources to do this."

In January 2017, a poker-playing supercomputer built by HPE and researchers at Carnegie Mellon University in the US handily beat four top professional players. Its specialty—out-bluffing humans—represents a milestone in AI.

"The goal is to translate this into an AI system that can handle and make decisions based on incomplete information," said Goh. "This is very relevant because you need to be more self-sufficient the further out you go from Earth."

While the harsh conditions of space can complicate even the simplest of tasks, there is one thing that is easier than on Earth—keeping supercomputers cool. The Spaceborne Computer is chilled with a fluid-cooling system that simply circulates through the frigid shadow side of the ISS.

"One could say it's the greenest supercomputer on—I almost said on Earth and I would have been wrong," laughs Goh, who adds that he is clearly on new ground. "One could say it's the greenest supercomputer built by humankind." ■

> " ON EARTH, YOU CAN AMASS RESOURCES TO GET INFORMATION TO MAKE DECISIONS. BUT IN SPACE, YOU HAVE LESS TIME AND RESOURCES TO DO THIS."

Dr. Goh Eng Lim
Vice President & Chief Technology Officer, HPC and AI Hewlett Packard Enterprise

NEVER TOO YOUNG TO START

Bringing supercomputing into the classroom

Age is just a number when using supercomputers to solve complex scientific problems.

By **Jeremy Chan**

Alibaba founder Jack Ma once said, "Young people will have the seeds you bury in their minds, and when they grow up, they will change the world." What Ma was trying to convey in his quote was the importance of providing young people with the relevant knowledge and skills to succeed as adults.

While textbooks and lectures are the de facto teaching method used in schools, educators now believe that experiential learning, where students are exposed to and challenged with real-world problems, can greatly enrich the education process. How educators plant the seed is arguably as important as what seed is being planted.

With this philosophy in mind, Mr. Harman Johll, head of department (research) at the National Junior College (NJC) in Singapore, is constantly on the lookout to add value to his college's research programs. Hence, when the National Supercomputing Centre's (NSCC) ASPIRE 1 supercomputer came online in 2016, he leapt at the opportunity to remotely tap on its 30,000 cores and one petaFLOPS computing prowess to take his students to the bleeding edge of scientific discovery.

"Leveraging the supercomputer at the NSCC, we're training our students to be both computational and experimental scientists," said Johll in an interview with *Supercomputing Asia*. "When our students go out to work, even when they're working as experimentalists, they ought to have the capacity to model and run their own base calculations using computational methods."

SPEEDY SCIENCE AT SCALE

Before ASPIRE 1 was built, Johll had cobbled together a computing cluster with the processing power to model, analyze and predict the properties of various materials. His rudimentary setup consisted of a handful of desktop computers linked by an ethernet network.

Impressively, with this workstation, he and his students were able to publish research in peer-reviewed journals such as *Scientific Reports* and the *Journal of Applied Physics*. His students also presented their work at professional-level conferences, among them, the International Conference on Materials for Advanced Technologies.

Still, Johll felt a lingering discontent. With the resources at hand, he and his students were limited to calculations on relatively small systems, or had to reduce the scope of a project in order to use calibrated or benchmark parameters that ensured meaningful analyses. In addition, some computations took more than a week to complete. Ambition and curiosity were bumping against the glass ceiling of hardware, and Johll was searching for a better way forward. Then, the ASPIRE 1 supercomputer came along.

"With supercomputing, we are no longer constrained by the size and complexity of the system we're investigating," said Johll. One of Johll's students, 17-year-old Ms. Pang Wen Ni, agreed, explaining how she can now work with 72 carbon atoms instead of just 32 when she runs analyses on the thermodynamics of fluorine displacing carbon atoms from graphene.

"The supercomputer allowed us to calculate the energies of the products and reactants, and get a clearer idea of which possible mechanisms the process could take. Without the supercomputer, it would be impossible for us to get such data," she said. Her schoolmate, 16-year-old Mr. Kee Jing Wen also found the supercomputer indispensable for calculating how hydrogen interacts with borophene, a two-dimensional analogue of graphene consisting of 36 boron atoms.

Johll further noted that such calculations, which used to take a week or more to complete, now only require two to three days. ASPIRE 1 has thus enabled a quantum leap in the speed and scale of computational research at NJC.

SUPERCOMPUTING, DEMOCRATIZED

Johll's students appear unfazed by the fact that they are grappling with complex scientific questions, using a tool thousands of times more powerful and sophisticated than anything most adults would have encountered.

Seated in front of their laptops in the NJC computer lab, Pang and Kee can often be found poring over lines of programming code displayed beside spheres suspended in three-dimensional space. They expertly call up software functions and toggle experiment parameters, as if they are merely playing a computer game.

"At first, I had to learn many different commands—even scrolling had its own commands, and that took a bit of getting used to," said Kee. "But over time, navigation within the system became second nature to me. It was easy to set up multiple runs and export data into my computer."

Johll commented that he only steps in occasionally to help troubleshoot technical errors and tweak the programming code. However, more time is spent contemplating the rigor of the scientific questions being asked and performing data analysis, rather than figuring out how to use the supercomputer.

With teenagers like Kee readily professing their proficiency with supercomputing, it would appear that the technology has truly broken out of its ivory tower. Anyone with an inquisitive mind can now tap on high performance computing to answer important scientific questions.

"Most people are quite surprised that a secondary school student is using a supercomputer to do research. Many assume that only full-time scientists or university students have access to such technology," Kee quipped. "I consider myself privileged to be able to use ASPIRE 1, and I'm grateful for the help and support from my mentors. It certainly makes my learning journey more fun and enjoyable!"

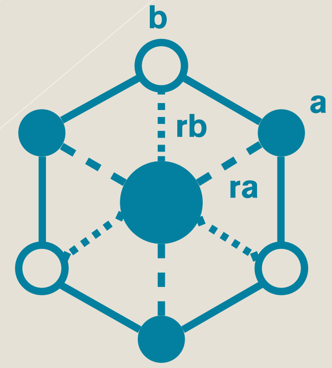

Modelling the interactions of atoms and molecules within novel materials is computationally intensive. ASPIRE 1 allows more complete and accurate simulations and measurements of these interactions to be performed.

COMPETE. COLLABORATE. CONNECT.

Highlights from SC17 in Denver, Colorado

Every November, the high performance computing community convenes in North America to compete, collaborate and, most importantly, connect with each other. The International Conference for High Performance Computing, Networking, Storage and Analysis—better known as SuperComputing or SC—was held in Denver, Colorado from November 12–17, 2017.

1


2


3

4

5

6


7

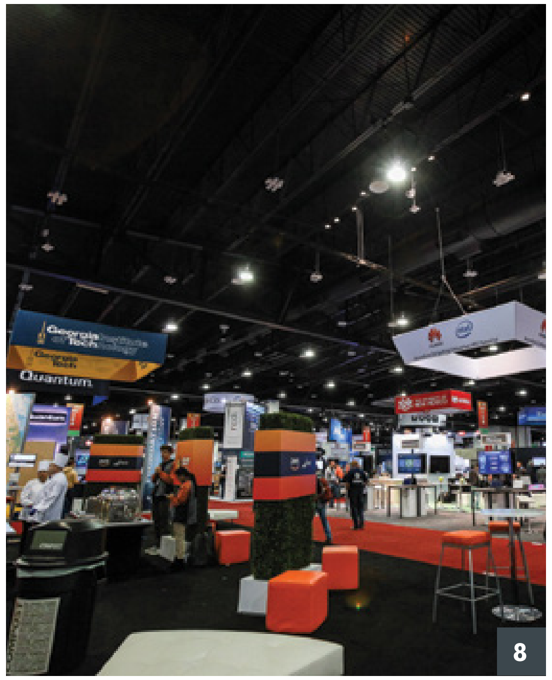
8

1. Welcome to Denver!
2. SC17 was attended by nearly 12,000 people, including scientists, developers and vendors.
3. Dr. Fu Haohuan, deputy director of the National Supercomputing Center, Wuxi, China, discussing the design philosophy of Sunway TaihuLight.
4. The 2017 ACM Gordon Bell Prize was awarded to a team from China that successfully ran a high frequency simulation of the 1976 Tangshan earthquake on the Sunway TaihuLight supercomputer.
5. The National Supercomputing Centre team representing Singapore on the exhibition floor.
6. SC17 chairman Dr. Bernd Mohr feeling relaxed as the conference gets off to a smooth start.
7. Professor Gordon Bell was invited to share his reflections on 30 years of the eponymous ACM Gordon Bell prize and its impact on high performance computing.
8. Industry and academic exhibitors seized the opportunity to showcase their latest products and projects.

Credits:

Photos 1, 2, 3:
NSCC Singapore

Photos 4, 5, 6, 7, 8:
SC Photography

9

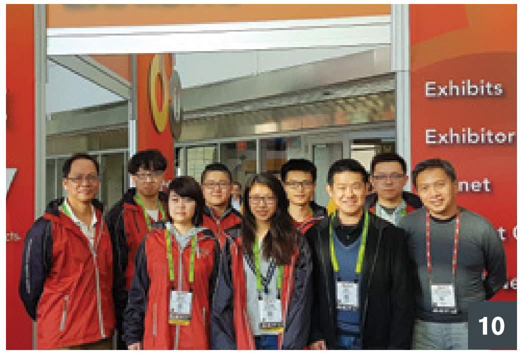

10

11

12

13

14

15

16

9. A total of 16 teams took part in SC17's student cluster competition.
10. After a close fight, the team from Nanyang Technological University, Singapore beat crowd favorites Tsinghua University to go home champions.
11. Associate Professor Judy Qiu of Indiana University discussing her group's research on data management to support high performance computing in the Apache Big Data Stack.
12. Professor Theresa L. Windus of Iowa State University giving an invited talk on the challenge of computational chemistry—scaling the very small to very large computers.
13. An early career speed-mentoring session in progress.
14. At the annual 'Dinner with Interesting People' hosted by Students@SC.
15. Professor Jack Dongarra takes to the stage during a special invited talk on the past, present and future of the TOP500 list. Check out our interview with him on p. 22!
16. Till we meet again in Dallas 2018!

Credits:

Photos 10,15: NSCC Singapore

Photos 9,11,12,13,14,16: SC Photography

Business Bytes

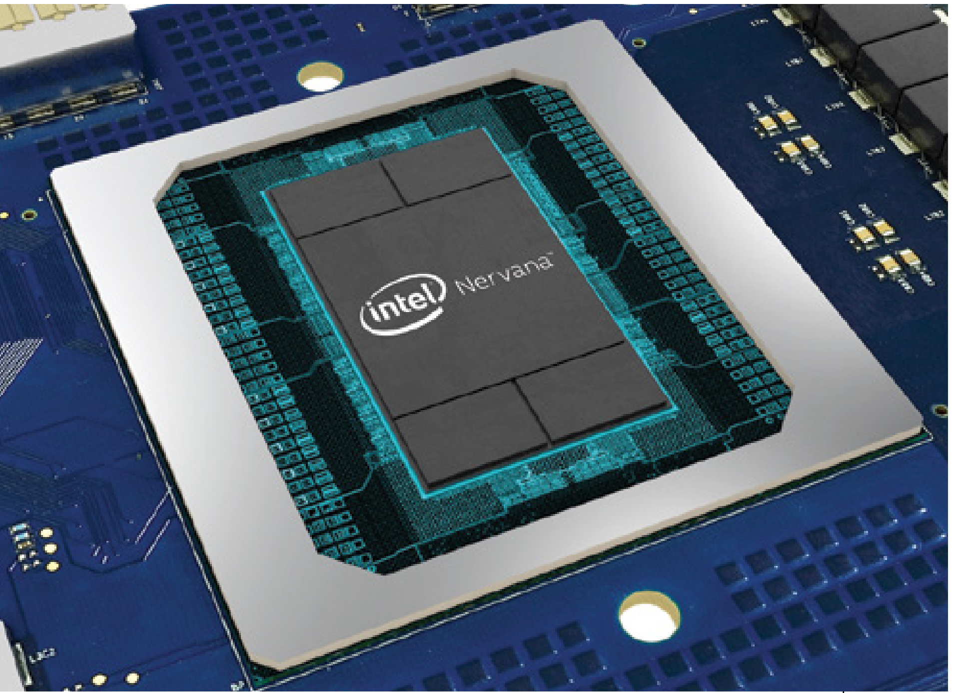

Intel's Nervana Neural Network Processor. Credit: Intel.

INTEL TO BEGIN SHIPPING NEURAL NETWORK CHIP

Intel has announced that their Nervana Neural Network Processor (NNP) will be commercially available by the end of 2017. Designed for deep learning applications, NNP is based on technology originally developed by Nervana, a company acquired by Intel in 2016 for US$350 million.

Intel's move into AI processors is part of an industry-wide shift, reflecting the growing demand for machine learning capabilities. NVIDIA, for example, has experienced a 56 percent year-on-year increase in revenue, largely driven by demand for their AI-centric V100 GPU accelerator.

"We have multiple generations of Intel Nervana NNP products in the pipeline that will deliver higher performance and enable new levels of scalability for AI models," Intel Corporation CEO Mr. Brian Krzanich wrote on the company's official blog. "This puts us on track to exceed the goal we set [in 2016] of achieving 100 times greater AI performance by 2020."

Facebook, which Krzanich said is in "close collaboration" with Intel on the NNP, is likely to be one of the first customers, using NNP as an alternative to Google's proprietary Tensor processing unit (TPU) for their deep learning workloads.

ASIAN CHIPMAKERS JOIN THE AI FRAY

Not to be left behind in the rapidly growing AI space, Fujitsu and Huawei have also announced their own versions of AI-specific microprocessors.

Fujitsu's offering, called the Deep Learning Unit (DLU), is purported to be ten times better than the competition in terms of performance per watt. The custom-built chip is designed to handle large-scale neural networks at a high performance and low power consumption. To achieve these two conflicting aims, Fujitsu has developed a new chip architecture, involving a few large cores and many small execution cores.

On top of the first generation DLU set to be launched in 2018, the company has plans for multiple generations of the chip, starting with a second generation chip that will have an embedded host CPU.

Huawei, on the other hand, has integrated its Kirin 970 processor in its Mate 10 smartphone released in October 2017. The integrated CPU and GPU features a neural processing unit that allows it to handle 2,000 images per minute. The Kirin 970 is said to have 25 times the performance and 50 times the energy efficiency of its predecessor, Kirin 960.

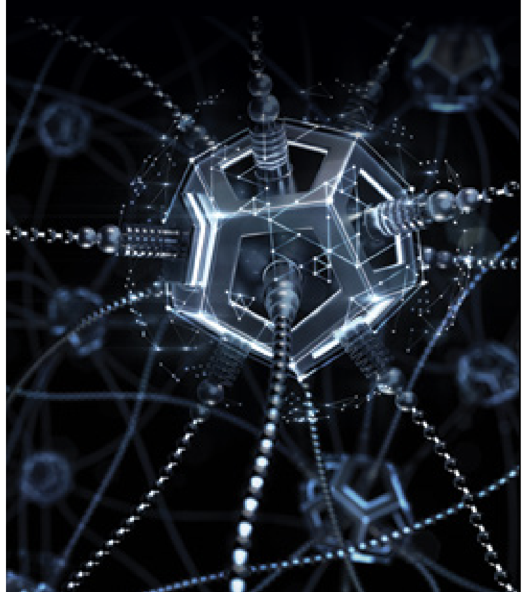

NEC IS BRINGING VECTORS BACK

First popularized by the likes of supercomputing pioneer Seymour Cray in the 1970s, the share of vector computers on the TOP500 list has declined over the years, in tandem with the rise of scalar computing powered by cheaper commodity microprocessors. With the launch of its SX-Aurora TSUBASA line, NEC Corporation hopes to bring vector-based computing back into the mainstream.

"The new product addresses the needs of scalar computational capability while still providing the efficiency of a vector architecture," said Mr. Kimihiko Fukuda, executive vice president, NEC Corporation, in a press release. "This is accomplished through a tightly integrated complete vector system in the form of a Vector Engine Card."

Unlike accelerators like GPUs that speed up a small portion of the code, the Aurora Vector Engine can execute complete applications, thereby removing the communication bottleneck between the host and the accelerator. It is also uniquely suited for applications with high memory demands, such as scientific and big data applications, as it has a memory bandwidth of 1.2 terabytes per second and a memory capacity of 48 gigabytes.

NEC's Aurora Vector Engine. Credit: NEC.

TOYOTA CHOOSES NVIDIA TO SUPPLY SELF-DRIVING CARS

NVIDIA's Drive PX platform has been selected to power self-driving cars for Toyota, the second largest automaker in the world. The Drive PX platform combines data from cameras, LiDAR, RADAR and other sensors, using AI to help the car understand and react to its environment.

In October 2017, NVIDIA announced its Drive PX Pegasus supercomputer, which the company claims will be able to operate at the Society of Automotive Engineer's 'level 5' classification. Level 0 cars are completely controlled by humans, while level 5 cars are considered fully autonomous, able to perform as well as a human driver even in extreme scenarios unlikely to be encountered by driverless cars, such as dirt roads.

Aside from Toyota, NVIDIA processors are also being used in Tesla, Mercedes and Audi cars. Competitor Intel, on the other hand, is collaborating with BMW to develop self-driving cars, and has acquired computer vision technology company Mobileye for US$15.3 billion.

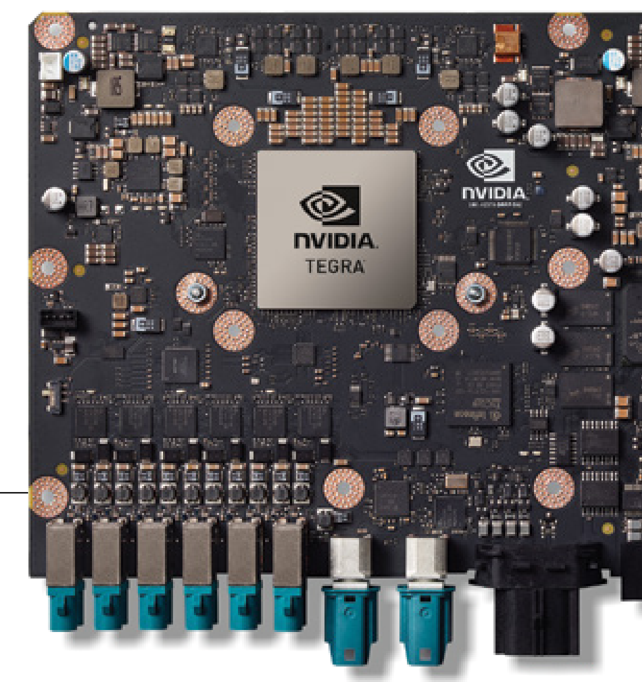

NVIDIA's Drive PX Pegasus. Credit: NVIDIA.

Super Snapshot

EVERYDAY WORDS THAT MEAN SOMETHING ELSE IN SUPERCOMPUTING

Like every field, supercomputing has its own jargon, unique lingo that specialists speak to get messages across faster. It can be confusing—not to mention difficult to 'Google'—so here's a glossary to help you figure it out!

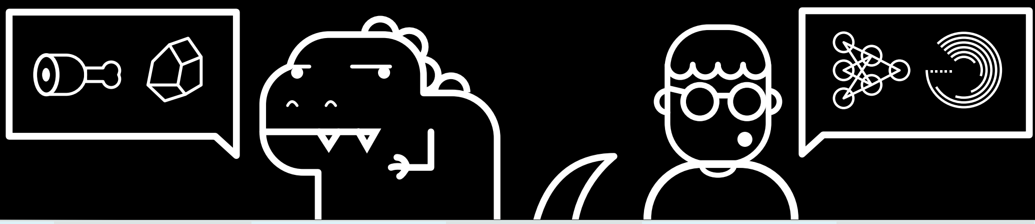

BLADE	CABINET	CORE	FLOP[S]	LIBRARY	SWITCH	THREAD
The cutting edge of a knife or other tool.	A committee of senior ministers in Commonwealth countries.	Muscles around the trunk, often trained by doing planks.	A total failure.	A collection of books or manuscripts.	A device that controls the flow of electricity.	A long, thin fiber used for sewing.
VERSUS						
A space-saving module that the computing, memory, power and networking components are physically mounted on.	Where multiple blades are stacked together and housed.	The basic unit of a processor. A multi-core processor contains two or more cores on a single chip.	Floating-point operations per second, a measure of how quickly a computer can operate.	A collection of data, pre-written codes or values that higher level programs can call upon.	Hardware that enables the multiple nodes of a parallel system to communicate. Also known as an interconnect.	The smallest unit of a program that can be managed independently by an operating system scheduler.

HPC in the Tropics

SINGAPORE
a little RED dot on the world map

Bali
Darwin
Malaysia
Philippines
Thailand
...and the list goes on

// Live and work in a tropical country; summer all year long

// An international business hub and a short flight away to many exotic tropical destinations

// Singapore has one of the most globalised economy in the world

// Immerse in a melting pot of international cultures and cuisines

// Experience a vibrant and multicultural society

Positions available:

ENGINEERS
We are recruiting a team of engineers to keep our supercomputing computing system in working order. We are seeking experts in the fields listed below. Candidates should possess a basic degree in computer science, IT or a related field from a recognised university as a pre-requisite. In addition, you should have sound working experience in the areas listed.

SYSTEMS
Candidates should be experienced in managing a compute cluster or a supercomputer, with a good level of proficiency in Linux as an administrator in a HPC environment. A firm understanding of job scheduling, preferably on PBS Pro, is essential.

STORAGE
Candidates should be proficient in GPFS and Lustre administration and tuning. A firm grasp of backup concepts is essential. Working knowledge of hierarchical storage systems is a definite plus.

CYBER SECURITY
Candidates will help to ensure that our cyber environment is safe from malicious attacks. The ideal candidate will be someone who has an overall grasp of the cyber security landscape rather than someone who merely runs cyber security analytics software.

CLOUD
Candidates will manage NSCC's cloud services. A sound working knowledge of OpenStack is essential.

NETWORKS
Candidates will maintain the supercomputer's external connectivity, as well as NSCC's connectivity to the rest of the world. You must be competent in, LAN, WAN, and structured cabling. As we run a mix of network devices from different suppliers, the ideal candidate should be someone with a sound understanding of the underlying concepts of network protocols and switching/routing. Certification is a welcomed bonus though not an absolute essential.

APPLICATIONS PROGRAMMER
Candidates will develop applications, either for NSCC's internal use or for enhancing our end-users' UNIX/Linux environment. Fresh graduates with relevant experience are welcome to apply.

HPC SOFTWARE ANALYSTS
We are hiring HPC Software Analysts in the various scientific domains to help our users run their codes more efficiently. In addition, we have a number of in-house technical projects in which you will be involved. You should posess at least a post-graduate degree (PhD preferred) in one of the computational sciences. You should also have substantial experience in parallel computing and programming optimising techniques. Experience in project management is a plus.

Please email your CV to careers@nscc.sg.

SG NSCC SINGAPORE
National Supercomputing Centre

For more information, please contact us at contact@nscc.sg or +65 6714 9450
1 Fusionopolis Way, Connexis South Tower, #17-01, Singapore 138632
www.nscc.sg

SCAsia Supercomputing 2018

Gathering the Best of HPC in Asia

26 - 29 MARCH 2018

RESORTS WORLD CONVENTION CENTRE, SINGAPORE

Organised by National Supercomputing Centre (NSCC) Singapore, Supercomputing Asia 2018 (SCA18) is the inaugural annual conference that encompasses an umbrella of notable supercomputing events. SCA18 will be held from 26 to 29 March 2018. The key objective is to promote a vibrant and relevant HPC ecosystem in Singapore. Delegates will be able to gain access to visionary insights from HPC luminaries thought leaders in academia and industry, optimum networking opportunities and the HPC community in Singapore.

WHAT'S IN STORE

INDUSTRY
HPC Technology Updates & Case Studies

SCIENTIFIC
Paper Presentations at Supercomputing Frontiers Asia

EDUCATION
Students' Experience Sharing & Cluster Computing Workshops

STRATEGIC SHOWCASE

Strategic verticals on the use of HPC
Eg: Health & Biomedical, Advanced Manufacturing & Engineering, Offshore & Marine, etc

CO-LOCATED HPC EVENTS

- **Asia-Pacific Advanced Network Meeting (APAN45)**
- **Conference on Next Generation Arithmetic (CoNGA)**
- **Singapore-Japan Joint Sessions**
- **Supercomputing Frontiers Asia (SCF Asia)**
- **Towards an Asia Pacific Research Platform (APRP)**

TUTORIALS

EXHIBITIONS

We expect to host many more delegates and prominent keynote speakers globally in the HPC space.

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

For sponsorships opportunities, please contact us at ***sponsorships@sc-asia.org***

For more information, please visit **SC-ASIA.ORG**