AsianScientist
Issue 01
January 2017

SUPERCOMPUTING ASIA

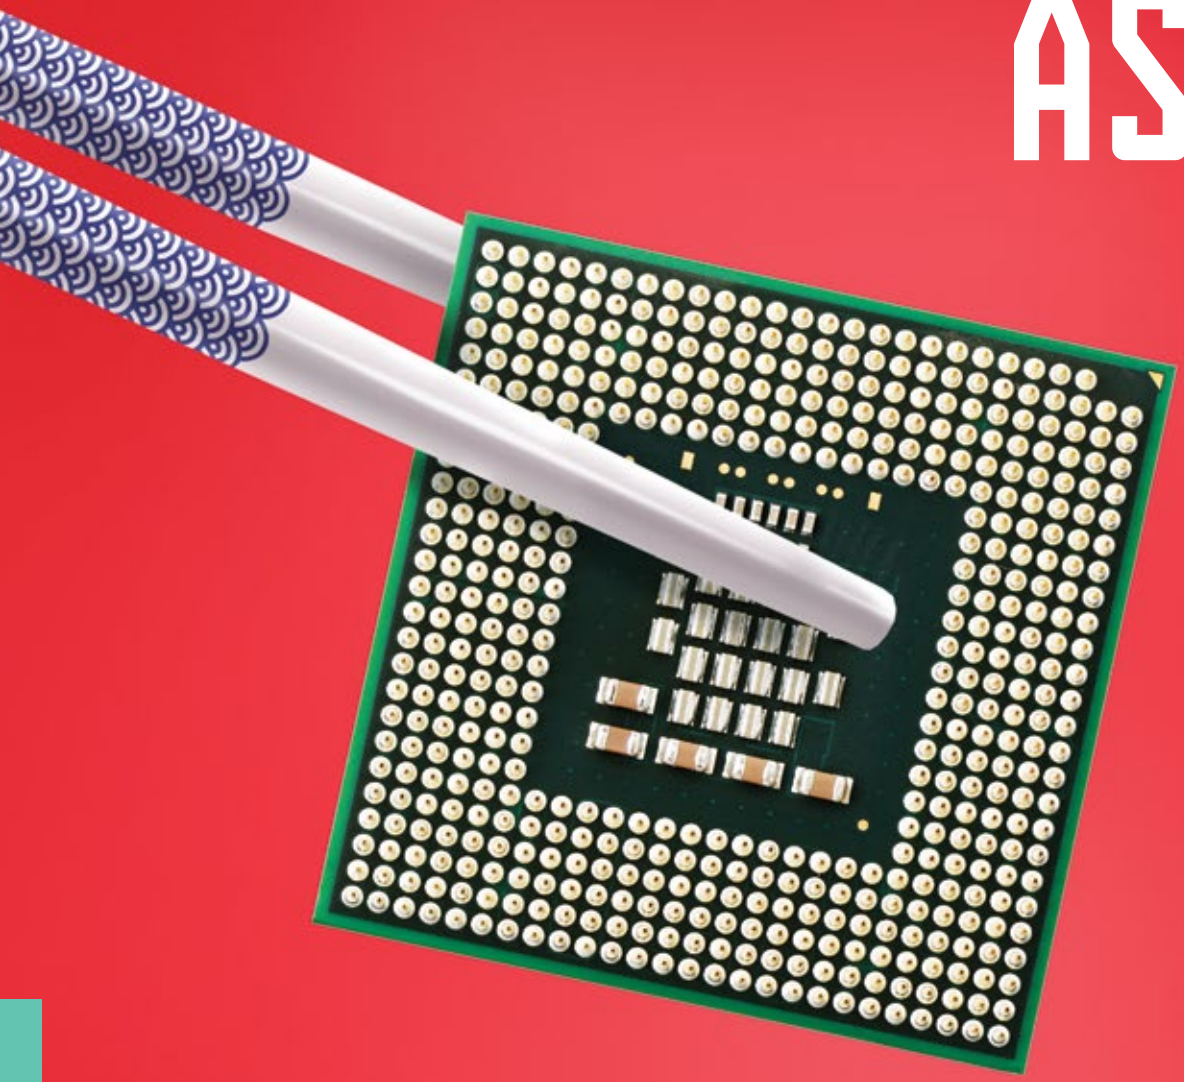

THE NEXT SUPERCOMPUTING SUPERPOWER

CHINESE TECHNOLOGY COMES OF AGE

ASIA'S SUPERCOMPUTING SCORECARD

POWERING THE LITTLE RED DOT

GREEN IS THE NEW BLACK

CONTENTS

Issue 01
January 2017

COVER STORY

FEATURES

EDITOR'S NOTE

In a rapidly evolving field like high performance computing, no one can afford to rest on their laurels. This is exemplified in the rise of China, which in a few short years has eclipsed the traditional powerhouses of the US and Japan. Just how did they develop the world's fastest supercomputer using home-grown chip technology? We examine this and other astonishing accomplishments of the Middle Kingdom in our cover story (*The Next Supercomputing Superpower*, p. 16).

But supercomputers are not just getting faster, they're getting greener too. Satoshi Matsuoka (*Green Is The New Black*, p. 34) shares with us how he made Tokyo Institute of Technology's Tsubame-KFC computer super-efficient by immersing the entire machine in oil. Our other interviewee in this issue needs no introduction—he is computing pioneer and legend Gordon Bell of the eponymous ACM Gordon Bell Prize (*An Exemplary Engineer*, p. 32).

Finally, we turn our attention to Singapore, a small nation that has an outsized supercomputing capability (*Powering The Little Red Dot*, p. 28). With its own petascale facility, Singapore holds its own against other Asian nations (*Asia's Supercomputing Scorecard*, p. 20).

Juliana Chan
Editor-in-Chief

SUPERCOMPUTING ASIA

EDITORIAL ADVISORY COMMITTEE
Prof. Tan Tin Wee
Nebojsa Novakovic
Dr. John Gustafson

EDITOR-IN-CHIEF
Prof. Juliana Chan

MANAGING EDITOR
Dr. Rebecca Tan

CONTRIBUTORS
Jerry Lim
Jon Lau
Cindy Lim
Nurfilzah Rohaidi
Jeremy Chan

SALES & MARKETING
Teo Yiling

DESIGN & LAYOUT
Lam Oi Keat

PUBLISHED BY
Asian Scientist Publishing Pte Ltd

DISTRIBUTED BY
Pansing Distribution Pte Ltd

HEAD OFFICE
Asian Scientist Publishing Pte Ltd
5 Toh Tuck Link, Singapore 596224
Tel: 65-6466-5775 Fax: 65-6467-7667
Website: www.asianscientist.com/subscribe
Facebook: www.facebook.com/asianscientist
Twitter: www.twitter.com/asianscientist

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

Digital Dispatch

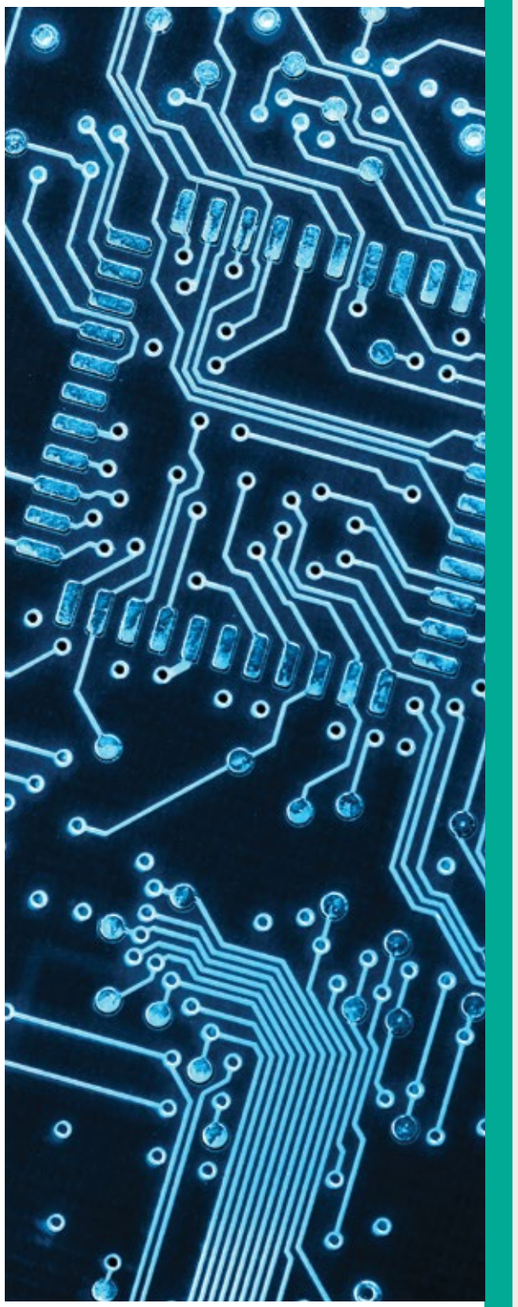

IBM'S WATSON DETECTS RARE LEUKEMIA IN MINUTES

IBM's Watson supercomputer has slashed the time taken to diagnose rare leukemias to just ten minutes—a feat considering that a human doctor would have taken two weeks to reach a similar diagnosis.

Doctors had initially diagnosed a Japanese patient with acute myeloid leukemia, a type of blood cancer. Soon enough, however, they realized that they were looking at a different type of leukemia, and looked to the cloud-based, artificial intelligence-powered supercomputer for a solution.

By cross-referencing and analyzing data from tens of millions of oncology papers, Watson detected over a thousand genetic mutations in the patient's DNA, filtering out the ones that were diagnostically important.

The researchers then concluded that the patient had a rare secondary leukemia caused by myelodysplastic syndromes, a group of diseases in which the bone marrow makes too few healthy blood cells.

"It might be an exaggeration to say artificial intelligence saved her life, but it surely gave us the data we needed in an extremely speedy fashion," Professor Arinobu Tojo told *The Japan Times*.

INDIAN INSTITUTE OF TECHNOLOGY GUWAHATI LAUNCHES SUPERCOMPUTING FACILITY

Param-Ishan, a new high performance computing system at the Indian Institute of Technology Guwahati (IIT-Guwahati), has recently been unveiled.

The supercomputer was installed in a partnership between the Centre for Development of Advanced Computing and IIT-Guwahati, providing a boost to the high performance computing community in India. Param-Ishan has a capacity of 300 terabytes and computing speed of 250 teraFLOPS, or 250 trillion operations per second.

Potential applications of the supercomputer lie in the fields of computational chemistry, computational fluid dynamics, and civil engineering structures. Notably, in a country that receives heavy rainfall during the monsoon seasons which inevitably leads to flooding and landslides, Param-Ishan could be useful in climate modelling and seismic data processing.

Professor Gautam Biswas, director of IIT-Guwahati, said that the supercomputer will not only augment the research initiatives in the institute, but also help in creating an ecosystem for attracting the right research talents, as reported by India's *NDTV*.

JAPAN'S SUPER-KAMIOKANDE DETECTOR GETS COMPUTING BOOST

Kamioka Observatory, part of the Institute for Cosmic Ray Research at the University of Tokyo, has put in an order to Fujitsu for an experiment-analysis computer system capable of handling the massive amount of neutrino observation data from its Super-Kamiokande detector.

Due to the delicate nature of neutrino observation, the detector must always be ready to reliably observe fleeting phenomena. As such, the computer system must not only be able to operate with stability all year round, it must also be able to perform high-speed analysis and process and store up to 500 GB of data per day.

Slated for operation in March 2017, the system will have a performance level about triple that of the existing system, as well as about triple the disk space and data transmission speed.

The system will accumulate observational data and analyze factors such as energy and the direction that neutrinos move. Analyses like these will hopefully shed light on the properties of neutrinos, or perhaps solve the mystery of how matter was created immediately after the Big Bang.

Digital Dispatch

INFINICORTEX INTERCONTINENTAL NETWORK LAUNCHED

The global supercomputer industry is faced with the challenge of ramping up high performance computing (HPC) systems to exascale capabilities while keeping the electrical power consumption for a system of that scale to less than 20 MW. One possible solution is to use distributed supercomputers to alleviate intense power requirements at any single location.

The InfiniCortex project led by Singapore's Agency for Science, Technology and Research (A*STAR)'s Computational Resource Centre (A*CRC) is a worldwide InfiniBand fabric that brings together several supercomputing facilities spanning four continents: Asia, Australia, Europe and North America.

According to Dr. Gabriel Noaje, senior computational scientist at A*CRC, the InfiniCortex high-bandwidth intercontinental network enables concurrent supercomputing across the globe, at speeds of up to 100 Gbps. Such an architectural approach could help build the next generation of HPC systems that solve complex problems through the aggregation and parallelization of many globally distributed supercomputers into a single hive-mind.

WHAT BUSINESS PROBLEMS CAN YOU SOLVE USING A SUPERCOMPUTER?

Since the launch of the National Supercomputing Centre Singapore, there have been a flurry of events to spread awareness of the new facility. One of these events was a hackathon that invited participants to "play with one of the most powerful supercomputers in Southeast Asia." The supercomputer in question has a speed of one petaFLOPS, capable of one quadrillion mathematical computations per second!

The hackathon, which took place from August 6–7, 2016 at the Fusionopolis research hub in Singapore, asked the question: What business problems can you solve using a supercomputer?

ModStore, the winning concept by a team of researchers from A*STAR's Institute of High Performance Computing, is a large-scale data science platform that allows users to learn, build and deploy data science models. From these models, users may extract business patterns, add value to services, and enhance business operations.

For their efforts, the winning team was presented with the hackathon's first prize of half a million CPU core hours, high-tech gadgets and cash prizes totaling over S$5,000.

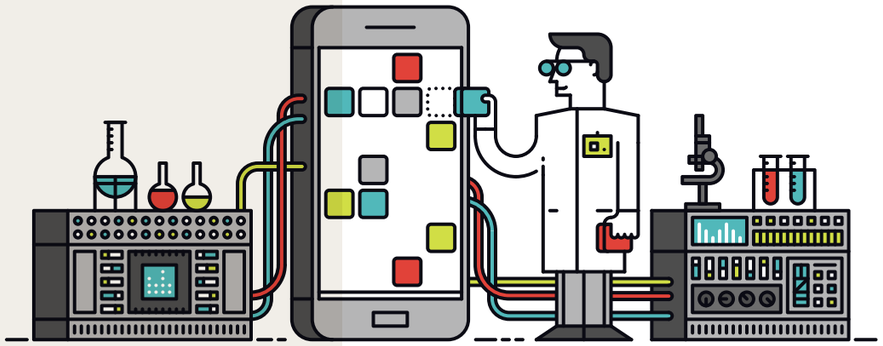

WHAT'S UP!

SUPERCOMPUTING FRONTIERS CONFERENCE BACK FOR THIRD YEAR

WHAT

Supercomputing Frontiers 2017

WHEN

March 13–16, 2017

WHERE

Level 4 Matrix Building, Biopolis 30 Biopolis Street, Singapore 138671

Singapore's annual international supercomputing conference, Supercomputing Frontiers 2017, will be back this coming March. Organized by the National Supercomputing Centre Singapore, the conference will provide a platform for thought leaders from academia and industry to explore global trends and innovations in high performance computing (HPC).

Now in its third year, the conference will include sessions on efforts to build exascale supercomputers; new, non-standard processor architectures; supercomputing and cryptography methods; integration of computing hardware, storage and networking in HPC and Big Data domains; and HPC applications in science, medicine and business.

Featuring an exciting lineup, distinguished keynote speakers for the conference include Gordon Bell, researcher emeritus at Microsoft Research Silicon Valley Laboratory (see interview on pg. 32), Alessandro Curioni, Europe vice president and director of IBM Research—Zurich, Thom Dunning, director of the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign and Fu Haohuan, deputy director of the National Supercomputing Center, Wuxi. Selected papers for the conference will be published in the international journal, *Supercomputing Frontiers & Innovations.*

For more information, visit supercomputingfrontiers.com.

ASC STUDENT SUPERCOMPUTER CHALLENGE 2017 NOW CALLING FOR APPLICATION PROPOSALS

WHAT

ASC Student Supercomputer Challenge 2017

WHEN

April 2017

WHERE

National Supercomputing Center Wuxi, China

The world's largest supercomputing hackathon—made possible by the Asia Supercomputer Community (ASC)—is now accepting application proposals and team registrations. The ASC Student Supercomputer Challenge 2017, or ASC17, seeks to promote the cultivation of young high performance computing talent, and in doing so, push for greater innovation in supercomputing. University students from all over the world are welcome to participate.

The competition will be hosted by the National Supercomputing Center in Wuxi, China, home of the Sunway TaihuLight supercomputer, which was ranked the top supercomputer in the world in 2016.

Applications must meet the following requirements: the typical use of the application is for scientific research and industry operations; the application should include the source code, installation package, and license; Linux-based applications should be supported by MPI Parallel; and, if shortlisted, the application owner should be able to provide complete technical files.

To participate, fill out the application form which can be found at www.asc-events.org/ASC17/ and send it to techsupport@asc-events.org by January 31, 2017.

For more information, visit www.asc-events.org/ASC17/

10 AWESOME REASONS TO LOVE SUPERCOMPUTERS

A layman's guide

From their sheer scale to the multitude of things they can do, there are many reasons to love supercomputers. Here are just ten of them!

By **Jeremy Chan**

The ongoing quest by scientists to develop better and faster supercomputers could be described as the human brain's attempt at creating its equal. Today, supercomputers are already helping researchers gain insight into the complicated workings of our world and ourselves. Check out these ten features that put the "super" in supercomputer, and decide for yourself if the human brain has succeeded in rendering itself obsolete.

1

THEY'RE SUPER FAST

Supercomputers can handle much more information at once than the average desktop. Consider a plane-load of incoming tourists at an airport customs checkpoint. Only one counter is open, so only one passport is cleared at a time. This is serial processing, which takes place in normal computers.

A supercomputer, by contrast, is akin to having thousands of open counters at the customs checkpoint, allowing passengers to be split among the counters and cleared more rapidly. This is parallel processing, which puts the 'super' in supercomputing.

A desktop computer has a serial processing power of approximately 40,000,000,000 FLOPS (floating point operations per second—a standardized unit of measurement for processing power). The parallel processing power of the world's fastest supercomputer, China's Sunway TaihuLight, is 2.5 million times greater, at nearly 100 petaFLOPS. That's a string of fifteen zeroes after the number 100!

2

THEY'RE HOT

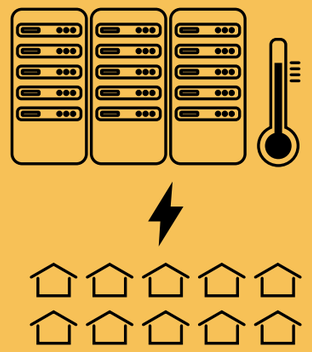

Like brains that (arguably) run on coffee, each central processing unit (CPU) needs electrical energy to crunch data. Unsurprisingly, a supercomputer with thousands of CPUs consumes megawatts of electrical power. As a rough estimate, the amount of electrical energy supplied to China's Tianhe-2 to keep it running daily is enough to power more than 10,000 homes, or a small city. All this electrical energy gets dissipated as heat, meaning you would break a sweat just by being close to a supercomputer.

To keep supercomputers from overheating, air conditioning or liquid coolant systems must be installed. All these additional features make supercomputing an energy-intensive and expensive activity, not to mention the impact on the environment. To reduce the environmental footprint of supercomputing, heated water from supercomputers is sometimes used to keep buildings warm. See our interview with green computing pioneer Satoshi Matsuoka on pg. 34 for more info!

3

THEY'RE BIG

In supercomputing, size matters. Your laptop is small because it contains only two to four central processing units (CPUs) which is sufficient for you to perform basic tasks such as surfing the internet or creating word documents. Supercomputers are made up of thousands of CPUs packed closely together in rows upon rows of cabinets with very fast networks interconnecting them. This allows a supercomputer to carry out much more complicated functions in the same amount of time or less.

To return to the airport analogy, if each customs counter takes up two square meters of space, to have a thousand of them set up would require proportionately more room. This is why supercomputers like the Tianhe-2 in China occupy a whopping 720 square meters of space. That's the size of one and a half NBA basketball courts!

LAPTOP

VS

SUPERCOMPUTER

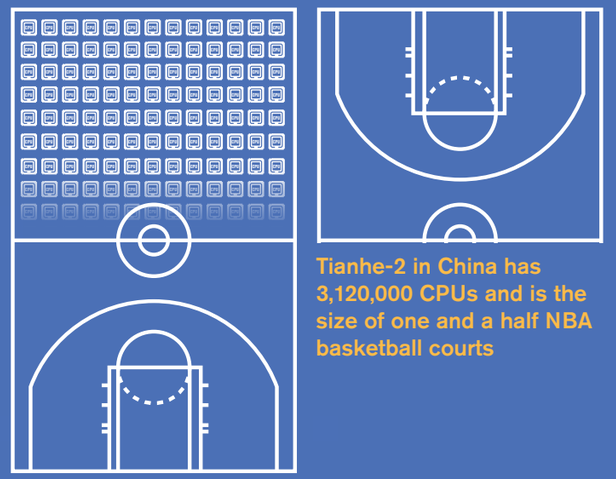

Tianhe-2 in China has 3,120,000 CPUs and is the size of one and a half NBA basketball courts

4

THEY CAN PREDICT THE WEATHER

Does weather forecasting sound suspiciously like fortune-telling, where someone looks into a crystal ball and predicts the future? Well, rest assured that meteorologists depend on something much more reliable than a crystal ball to make weather predictions. All around the world, supercomputers are hard at work processing real-time weather data, monitoring temperature, humidity, wind speed and cloud cover around the world to find weather patterns.

Making sense of all this information for a weather forecast is practically impossible for the human mind, or it could take so long that one might as well look out the window for an answer. For a supercomputer, however, it's all in a day's work; maybe less. The next time you get a storm warning and reach for your umbrella, remember that you have a supercomputer to thank for keeping you dry.

5

THEY CAN PREDICT EARTHQUAKES AND TSUNAMIS

The 2011 Great East Japan Earthquake that caused the Fukushima Daichi nuclear power plant meltdown exceeded all expectations by geologists in terms of the extent of damage and loss of life. As simulations for tsunami flooding require very complex mathematics, they take several days when run on a regular computer—much too long to be useful for real-time analysis and activation of early warning systems.

Using the K supercomputer that was fed with data from the Great East Japan Earthquake, a Japanese research group has dramatically winnowed down the time required for such simulations to just two minutes. It does so with remarkable accuracy when used to model real-life flood scenarios. This means that earthquakes and tsunami flooding can be predicted and evacuation notices sent out way ahead of time, long before the waves hit the shore.

6

THEY CAN BEAT GRANDMASTERS AT THEIR OWN GAME

In the year 1996, chess grandmaster Garry Kasparov faced off against Deep Blue, a supercomputer which could compute more than 200 million chess moves per second. The man fared admirably against the machine, losing in only one of six matches. Yet, that match was notable because never before had a supercomputer defeated a grandmaster in chess.

More recently, Google's supercomputer, AlphaGo, took on Lee Se-Dol, a South Korean Go grandmaster. The objective of Go is for one player, using black pieces, to surround another player's white pieces on a gridded board, and vice versa. The maximum number of possible moves for Go has been estimated to exceed the total number of atoms in the visible universe, making it an incredibly complex game. Lee lost three consecutive matches to AlphaGo before clawing back one victory in the fourth match and surrendering to AlphaGo in the fifth.

7

THEY'RE EARNING MILLIONS

Financial markets are a fickle lot. Market scenarios are constantly changing as new economic information gets released and all the market players react, often haphazardly. Although certain market movements might be predicted based on experience or intuition, not all situations can be manually sampled for risk and potential outcome.

A supercomputer capable of simulating the thousands of market transactions that might occur in response to new information would therefore be an extremely valuable tool for financial institutions. Indeed, J.P. Morgan has already enlisted the help of a Maxeler dataflow supercomputer in its fixed income trading operations. The 128 teraFLOPS processor allows complex trading analysis to be computed within minutes as compared to hours when performed by a human employee.

8

THEY DRIVE COOL CARS

Self-driving cars, also known as autonomous vehicles, have been touted as the future mode of transportation. But before they can be deployed, their safety on roads must be rigorously tested. Just think of how a human driver needs to react appropriately to other vehicles, pedestrians, traffic lights and weather conditions. An autonomous car must be able to do the same, with sensors as its 'eyes' and a powerful processor for its 'brain.'

Recently, NVIDIA, a company better known for its high performance graphics cards, has collaborated with the automotive industry to provide a supercomputer small enough to fit under the bonnet. The Drive PX2 is that supercomputer, packing the computing prowess of 150 MacBook Pros with the capacity to perform 24 trillion operations per second into the size of a lunchbox.

9

THEY CAN HELP DISCOVER NEW DRUGS

Most drugs used to treat human diseases take effect by binding to biological molecules called proteins in the body. A drug that fits snugly into 'pockets' on a protein's three-dimensional structure is deemed desirable. To design such a drug, however, the target protein's structure must first be known or predicted, and mapping the shape of a protein is a highly complex task. This is where computational modeling of protein structure becomes immensely helpful.

By performing extensive calculations on energy states and motion, D.E. Shaw Research's Anton supercomputer was able to accurately predict the three-dimensional shape and behavior of a protein. This opens the possibility for scientists to design better and more effective drugs with fewer side effects.

10

THEY CAN DIAGNOSE AND RECOMMEND TREATMENT FOR DISEASES

Diseases such as cancer are complicated, requiring customized therapy. For example, leukemia refers to a group of blood cancers that respond very differently to available treatments. Using genetic data from a patient, doctors can make better treatment decisions, but not before sieving through thousands of mutations to identify which specific leukemia subtype the patient suffers from.

Enter Watson, IBM's supercomputer which can compare a patient's genetic data against medical databases to make diagnoses. It accomplishes in ten minutes what a human doctor would have taken two weeks to achieve manually. Being an internet-based system, Watson can extend its diagnostic services beyond the US where its hardware is located. It has since been deployed in Asia, assisting doctors in Thailand and India. Move aside, Doctor House. Watson just stepped in.

THE NEXT SUPERCOMPUTING SUPERPOWER

CHINESE TECHNOLOGY COMES OF AGE

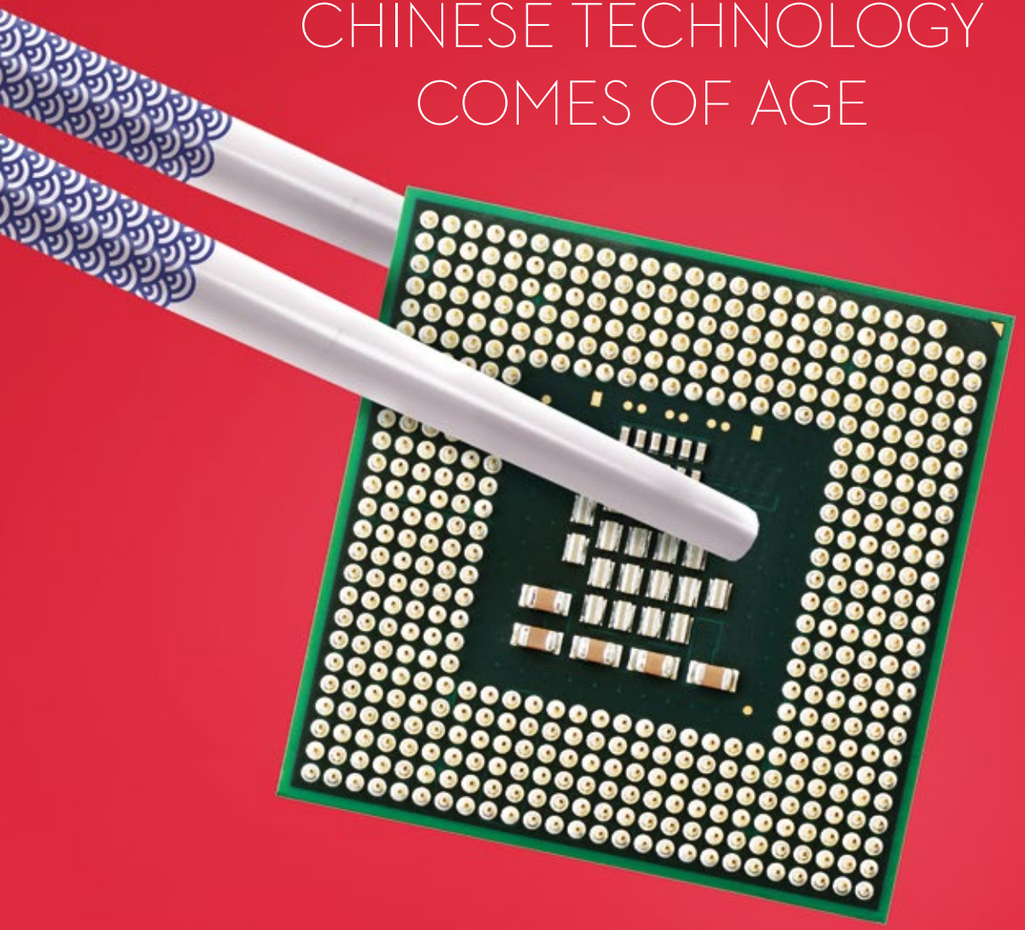

The most powerful supercomputer in the world uses China's home-grown technology. Could the country also be the first to build an exascale computer?

By **Rebecca Tan**

Which country has the most powerful supercomputer in the world? If you had asked that question when the first TOP500 list was published in June 1993, the answer would be the United States.

Back then, CM-5 at the Los Alamos National Laboratory claimed the crown of fastest supercomputer in the world, clocking in at 60 gigaFLOPS, or 60 billion floating point operations per second. A decade later, it would be Japan's turn to flex its supercomputing muscle, with NEC's Earth-Simulator topping the June 2003 ranking with a speed of 36 teraFLOPS.

But ask that question today and the answer would undoubtedly be: China. Chinese supercomputers have comfortably sat on the top of the biannual list since June 2013, when Tianhe-2 achieved a Linpack benchmark speed of 34 petaFLOPS—a staggering 1,000 times faster than leading machines in the preceding decade. In fact, Tianhe-2's number one spot would remain unchallenged for six consecutive TOP500 lists.

The accomplishments of Tianhe-2, meaning 'Milky Way' in Mandarin, are all the more impressive considering that China did not even appear on the TOP500 list until as recently as 2001. China's rise since then can only be described as meteoric, with Tianhe-1 breaking into the top ten in 2009, swiftly followed by Tianhe-1A and Nebulae taking the first and third spots in 2010.

China's recent dominance on the TOP500 list has not gone unnoticed, said Jack Dongarra, a professor at the University of Tennessee and compiler of the TOP500 list since its inception. "It's part of an overall trend in China; they had zero systems in 2001 and today they surpass the US," he told *Supercomputing Asia*. "No other nation has seen such rapid growth."

SEEING CHINA IN A WHOLE NEW LIGHT

Despite the growing recognition of China's capabilities in a field traditionally dominated by the US and Japan, Tianhe-2's number one ranking in 2013 still came as a surprise to many industry watchers as it was only expected to be deployed two years later. Surprise gradually gave way to alarm, eventually culminating in the 2015 US government ban on the sales of Intel, NVIDIA and AMD microchips to China on fears that the technology would be used for nuclear applications.

Ever since nuclear explosive testing was outlawed in 1996, simulations run on supercomputers have become more important for nuclear research. Furthermore, supercomputers also play an increasingly significant role in national cybersecurity efforts. Although China has repeatedly insisted that its supercomputers are being used for peaceful research, the fact that the second most powerful supercomputer in the world happens to be built by the National University of Defense Technology, has not helped to allay concerns.

In any case, the lack of access to US microchips does not appear to have slowed China down. "The ban has had an insignificant impact," said Deng Yuefan, a professor at Stony Brook University and a Mount Tai scholar at the National Supercomputer Center of China in Jinan, in an interview with *Supercomputing Asia*.

"I can list a few effects: Firstly, it has caused Intel to lose business in China. Secondly, it has somewhat delayed a planned upgrade of Tianhe-2. And lastly, it has hastened the development of the many-core 64-bit RISC 1.45 GHz SW26010 chips, which were designed by Shanghai's National High Performance Integrated Circuit Design Center."

Photo credit: Science China Press

Armed with the new Shenwei SW26010 chips, a new supercomputer has surged to the top of the TOP500 list: Sunway TaihuLight. It is nearly three times as fast as Tianhe-2, being benchmarked in Linpack as being able to perform 93 quadrillion calculations each second (93 petaFLOPS). To put this achievement in context, modern desktop PCs are already more powerful than the top ranked supercomputers from the early 1990s, and 93 petaFLOPS is roughly equivalent to the processing power of a million such PCs.

"The Sunway TaihuLight is simply a *tour de force*," Dongarra added.

BEYOND THE NEED FOR SPEED

However, simply owning the fastest supercomputer does not make China a supercomputing superpower by default: its might has to be demonstrated in how well all that power is put to use in applications. Indeed, Tianhe-2 has previously been criticized as a 'stunt machine' that has impressive specifications but lacks the software that makes it easy for researchers to use, on top of being too expensive to run.

"The widespread obsession with supercomputer hardware is understandable. No one who helps administer the TOP500 Supercomputer Sites project, as I do, can claim to be immune to it," Dongarra mused. "But when it comes to advancing the cause of computational modeling and simulation as a new part of the scientific method, there is no doubt that the complex software ecosystem must take centerstage. Software is the potential blind spot for any supercomputer architecture."

"Focusing on developing a supercomputer that ranks high on the TOP500 list but is only good for solving linear algebra problems is not a good use of funds," agreed Deng. "China has recognized the value of software development and is investing accordingly."

As evidence of this investment, three of the six finalists for the prestigious Association for Computing Machinery (ACM) Gordon Bell Prize 2016 were teams that used Sunway TaihuLight. One of the teams, a 12-man group that included Associate Professor Fu Haohuan of Tsinghua University in Beijing, was announced the winner at the 2016 International Conference for High Performance Computing, Networking, Storage and Analysis held in Salt Lake City, Utah in November 2016.

Aside from its raw power, perhaps what makes Sunway TaihuLight so astonishing is that it is not only twice as fast as its predecessor but also three times as efficient. This benefit is derived from relying on modern, China-controlled instruction set architecture instead of the antiquated complex X86, as well as an elegant platform and cooling system.

"On the one hand, TaihuLight uses a fully automated, closed-coupled liquid cooling unit. On the other hand, the core operating frequency for Shenwei SW26010 proccessors is low, at just 1.5GHz," explained Fu, who is also the deputy director of the National Supercomputing Center at Wuxi where Sunway TaihuLight is housed.

"These factors greatly reduce energy consumption, giving Sunway TaihuLight an efficiency of 6.05 gigaFLOPS/W, which is comparable to other top international systems."

2020 VISION

As impressive as it may be, Sunway TaihuLight is just the tip of the iceberg of China's supercomputing ambitions. "High performance computing is recognized as a powerful engine for growth and the key to China's competitiveness in science and technology," Fu shared with *Supercomputing Asia*. "At the Wuxi National Supercomputing Center, we will make use of the resources provided by Jiangsu province, combined with funding from the 13th Five-Year Plan, to build a globally influential technology innovation center."

One important area China hopes to take a lead in is building the first exaFLOPS-scale supercomputer. In May 2016, the central government announced that its 13th Five-Year Plan includes funding for a program to develop the country's first exascale supercomputers—Jinan University's system, which is based on a similar ShenWei architecture as TaihuLight, as well as the ARM-based Tianhe-3—which are both expected to be ready by 2020.

In the race to exascale, China has serious competition in the form of old rivals Japan and the US. Japan was in fact the first to announce a comprehensive plan to build an exascale computer, launching Flagship 2020 in 2014. Meanwhile, the US is currently working on two computers that could topple Sunway TaihuLight in the near future: IBM's 200 petaFLOPS Summit supercomputer at the Oak Ridge National Laboratory and the 300 petaFLOPS Sierra at the Lawrence Livermore National Laboratory. In addition, the US Department of Energy has plans for two exascale machines by 2023, according to the Exascale Computing Plan which was made public in April 2016.

"I am quite optimistic that China is on track to reach exascale by 2020 for the following reasons: There is still more than three years to go; China has already completed a 125 petaFLOPS system; it is able and willing to spend a lot of money; and the ranking is based on Linpack, which China is already strong in," Deng said.

"China is big and supercomputing problems are expensive to solve, but its supercomputing communities are highly collaborative. We will require the cumulative efforts of future generations if we are to develop the necessary central processing units, complete software systems and application packages for these exascale computers to become a reality."

ASIA'S SUPERCOMPUTING SCORECARD

Pokémon Go style

Asia is home to both supercomputing heavyweights and rising stars. Here are six of the region's best machines according to the latest TOP500 list, reimagined as Pokémon cards.

By **Teo Yiling**

Every June and November, the global supercomputing community awaits the release of the TOP500 list with bated breath. A ranking of the world's most powerful supercomputers, the TOP500 list is the high performance computing version of the Fortune 500 list, giving an indication of the capabilities of movers and shakers in the field.

Since the TOP500 list was first launched in 1993, the US and Japan have jostled for top spot on the continually evolving rankings. In recent years, however, China has begun to make its presence felt, having held the number one spot for three years in a row with a commanding lead that seems unlikely to be lost in the near future (see our special feature on China on pg. 16).

But China is not alone in its drive to capitalize on the power of supercomputing. Elsewhere across Asia, countries are realizing how supercomputers could give them a technological edge and have quickly developed impressive capabilities of their own.

Inspired by Pokémon Go, a wildly popular augmented reality game developed by Niantic, we show six of Asia's best supercomputers as Pokémon cards. These six supercomputers represent the very best science and technology that each country has to offer—at least for now. Supercomputers have gone from gigaFLOPS in the 1990s to teraFLOPS in the 2000s and are now in the petaFLOPS range, roughly improving by almost three orders of magnitude per decade. As 2020 approaches, watch this space for Asia's first exaFLOPS machines!

#1

Sunway TaihuLight

CHINA

RANK 1/500

No chips from Intel? No problem. Sunway TaihuLight runs on domestically-developed ShenWei SW26010 processors, which pack in a whopping 260 cores on each chip. This makes Sunway TaihuLight almost three times as fast as the next fastest supercomputer in the world: Tianhe-2, also developed by China.

93 petaFLOPS	10,649,600	1.45 GHz
Top speed	Total cores	Processor speed

National Supercomputing Center in Wuxi

National Research Center of Parallel Computer Engineering and Technology (NRCPC)

Year launched 2016

#7

K Computer

JAPAN

RANK 7/500

Named after the Japanese word for 10 quadrillion, *kei* (京), K Computer made its first appearance on the TOP500 way back in 2011. Despite its relative old age compared to newer entrants on the list, K Computer is no slouch when it comes to performance, maintaining its number one ranking on the Graph500 rankings for three consecutive lists.

10.5 petaFLOPS	705,024	2 GHz
Top speed	Total cores	Processor speed

RIKEN Advanced Institute for Computational Science (AICS)

Fujitsu

Year launched 2011

#15

Shaheen II

SAUDI ARABIA

RANK 15/500

The Shaheen II supercomputer burst onto the scene in 2015, breaking into the top ten with its seventh place ranking. The highest ranked supercomputer in the Middle East, Shaheen II is almost 30 times faster than its predecessor, Shaheen, which is Persian for peregrine falcon. Major users include petrochemical companies such as Saudi Aramco and Saudi Basic Industries Corporation.

5.5 petaFLOPS	196,608	2.3 GHz
Top speed	Total cores	Processor speed

King Abdullah University of Science and Technology (KAUST)

Cray Inc.

Year launched
2015

#46 & #47

Miri & Nuri

SOUTH KOREA

RANK 46,47/500

The Korea Meteorological Administration's twin supercomputers Miri (미리) and Nuri (누리) were commissioned in 2015. Unlike other more general purpose supercomputers on the TOP500 list, Miri and Nuri are dedicated to climate research, and together form the largest operational numerical weather prediction system in the Asia-Pacific region.

2.4 petaFLOPS	69,600 x 2	2.6 GHz
Top speed	Total cores	Processor speed

Korea Meteorological Administration

Cray Inc.

Year launched
2015

#115

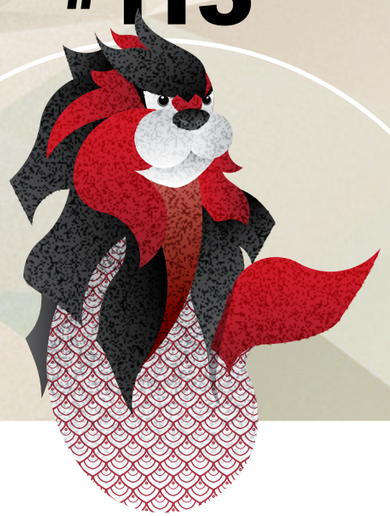

ASPIRE 1

SINGAPORE

RANK 115/500

Though much less populous than the other Asian countries featured on this list with a population of just over five million, Singapore nonetheless has petaFLOPS-sized ambitions. The National Supercomputing Centre Singapore's one petaFLOPS machine, Advanced Supercomputer for Petascale Innovation, Research and Enterprise or ASPIRE 1, made its debut on the TOP500 list in June 2016.

1 petaFLOPS	30,912	2.6 GHz
Top speed	Total cores	Processor speed

National Supercomputing Centre Singapore (NSCC)

Fujitsu

Year launched 2016

#133

SahasraT

INDIA

RANK 133/500

Sahasra means "a thousand arms," a nod to the multi-tasking that India's top supercomputer is expected to do. Indeed, SahasraT has already proven its mettle in both computational aerodynamics and astrophysics, having successfully simulated the landing sequence of a high lift wing and the overlapping of supernovae.

0.9 petaFLOPS	31,104	2.5 GHz
Top speed	Total cores	Processor speed

Supercomputer Education and Research Centre (SERC), Indian Institute of Science

Cray Inc.

Year launched 2015

POWERING THE LITTLE RED DOT

Shaping Southeast Asia's supercomputing landscape

Home to the only petaFLOPS-scale system in Southeast Asia, the National Supercomputing Centre Singapore handles large volumes of data generated by the small island nation.

By **Nurfilzah Rohaidi**

Though it has comparatively fewer citizens than the other countries in Southeast Asia, Singapore makes the most out of its most precious resource—its people. The small island nation houses some of the top universities in Asia and is internationally recognized as a research powerhouse that punches well above its weight. Correspondingly, Singapore's highly educated and tech-savvy population has a large appetite for high performance computing (HPC) to support its research efforts in both academia and industry.

Those needs have now been met in the National Supercomputing Centre Singapore (NSCC). As the first national center of petascale standard in Singapore and the wider region of Southeast Asia, NSCC is a plug-and-play facility that offers computing, multi-petabyte data storage and multi-gigabit speed network resources to enable users to solve scientific and technological problems.

"Here at NSCC, our aim is to use supercomputers to enable researchers to conduct research at the speed of thought."

Mr. Jon Lau,
Deputy Director
(Business Development)
National Supercomputing Centre Singapore

"In today's world, the availability of HPC is an indicator of the country's scientific and engineering prowess, and more importantly, an economic driver," said Mr. Jon Lau, deputy director (business development) at NSCC.

"Here at NSCC, our aim is to use supercomputers to enable researchers to conduct research at the speed of thought, leveraging technology to transform society and individual lives through research."

PREDICTING THE NEXT BIG ONE

For researchers like Sylvain Barbot, an assistant professor at Nanyang Technological University's Earth Observatory of Singapore (EOS), the NSCC facility is indeed an indispensable resource.

"HPC is critical to analyzing remote sensing data and modeling large geophysical datasets that help us understand our changing planet," said Barbot, who studies the physics of earthquakes and tectonic plate boundaries, seeking to describe how earthquakes start and why they cause so much destruction.

By offering flexible architecture that allows researchers to deal with the large datasets generated by remote sensing and field data, NSCC facilitates computationally demanding simulations and the intense data processing required to understand the complexity of the Earth's mechanical behavior, he added.

"In fact, HPC was central to many groundbreaking studies at EOS," Barbot told *Supercomputing Asia*. "For example, we have used HPC to study historical data on tsunamis in Southeast Asia to better understand current risk, and also how much shaking would occur in Singapore if an earthquake were to strike nearby Indonesia."

TRY BEFORE YOU BUY

The national center, which welcomes collaborations, has also enabled research that could potentially save the offshore and marine industry millions of dollars. Using NSCC facilities, researchers at the National University of Singapore's (NUS) Technology Centre for Offshore and Marine are working on using computational fluid dynamics (CFD) to allow companies to try out their ideas before committing to build expensive ships and rigs.

"One of the big problems in the oil and gas industry right now is that the price [of oil and gas] is quite low. People are now searching for innovative ways to produce oil and gas at a lower cost but still do it safely," Professor Allan Magee of the Centre for Offshore Research & Engineering at NUS, told *Supercomputing Asia*. "CFD simulations—coupled with physical testing—give companies design assurance that they haven't missed any of the important physics."

Noting that CFD is also widely used in fields such as airplane and building design, Magee said that the offshore and marine sector presents unique challenges that stretch HPC to the limit. "Waves are among the most complex flow problems that exist in nature because it is an unsteady problem and the ocean is gigantic. We need to resolve turbulence at the millimeter scale—imagine trying to do that over a domain which is kilometers-long. So, obviously, it takes a very powerful computer to do that," he explained.

Thankfully, NSCC is up to the task. "Only developed nations have similar capabilities," Magee added. "There's nothing in Southeast Asia that I'm aware of that rivals it."

Magee's research has also benefited from NSCC's high-bandwidth network, the InfiniCortex 'galaxy of supercomputers.' "It provides a backbone or magnet which connects us to all other supercomputers throughout the world so that even if the local computer is not able to solve the problem we can connect to others that might be able to," he said.

"This enables us to share data and computational power internationally, helping us to accomplish much more through synergy and collaboration."

DEMOCRATIZING HPC

Just as personal computers and mobile phones have improved connectivity across the board, supercomputing is expected to revolutionize entire industries and become widely used by researchers of all disciplines. Indeed, other researchers who have benefitted from NSCC's services come from diverse fields, ranging from molecular dynamics to microbiology and genomics.

"NSCC's long-term vision is to continually make HPC accessible to all, thereby democratizing HPC and building a healthy ecosystem of supercomputing users," Lau said. To encourage HPC adoption, NSCC conducts monthly training workshops that are customized to each user's needs. They also organize events to raise the awareness of what HPC can achieve, including the inaugural NSCC Hackathon held in August 2016, where participants used supercomputers to solve problems faced by local start-ups.

Ultimately, NSCC is not just a useful resource but a facility run by researchers for researchers.

"NSCC understands the requirements of us scientists," said Barbot. "We don't just work with them, we truly collaborate."

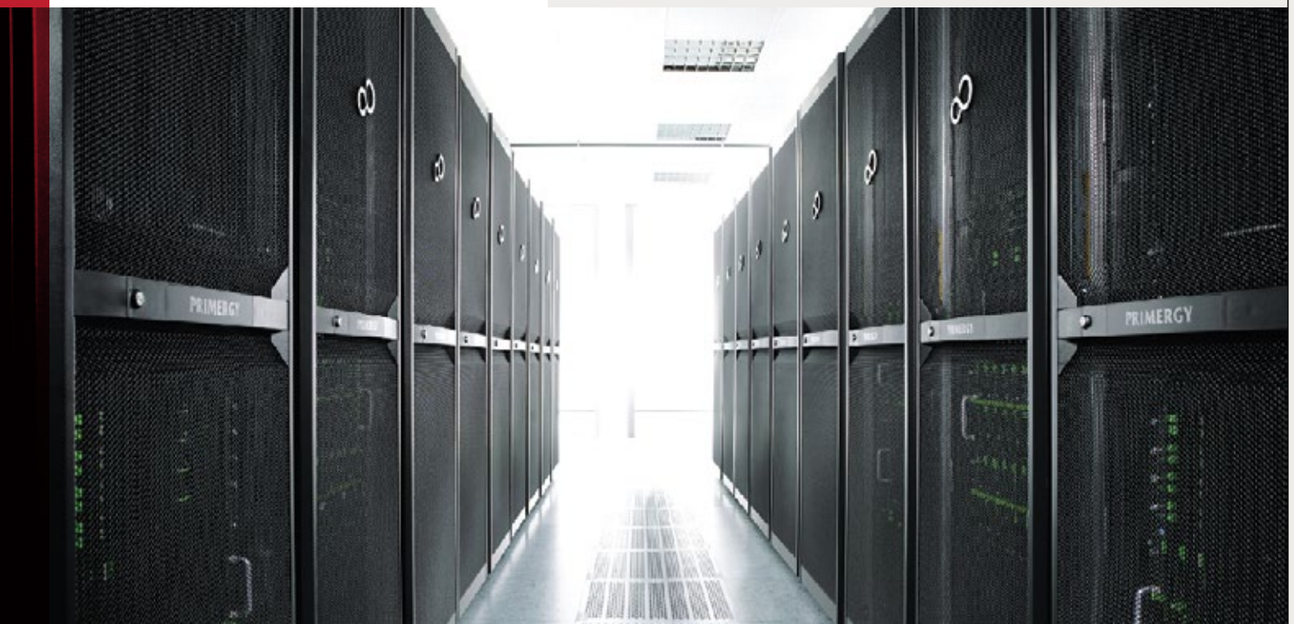

If you are interested in signing up for a trial beta account to use NSCC's HPC resources, please visit nscc.sg/contact-us or email contact@nscc.sg with the following information:

1. Name of user
2. Preferred User ID (limited to 8 characters)
3. Designation
4. Organization
5. Nationality of user
6. Purpose of using NSCC's supercomputer
7. Project name, if any
8. Number of CPU-cores hours required
9. Number of GPU-hours required
10. Storage required (GB)

AN EXEMPLARY ENGINEER

THE EXTRAORDINARY LIFE OF GORDON BELL

A *bona fide* pioneer of high performance computing, Professor Gordon Bell nonetheless sees himself first and foremost as an engineer.

By **Rebecca Tan**

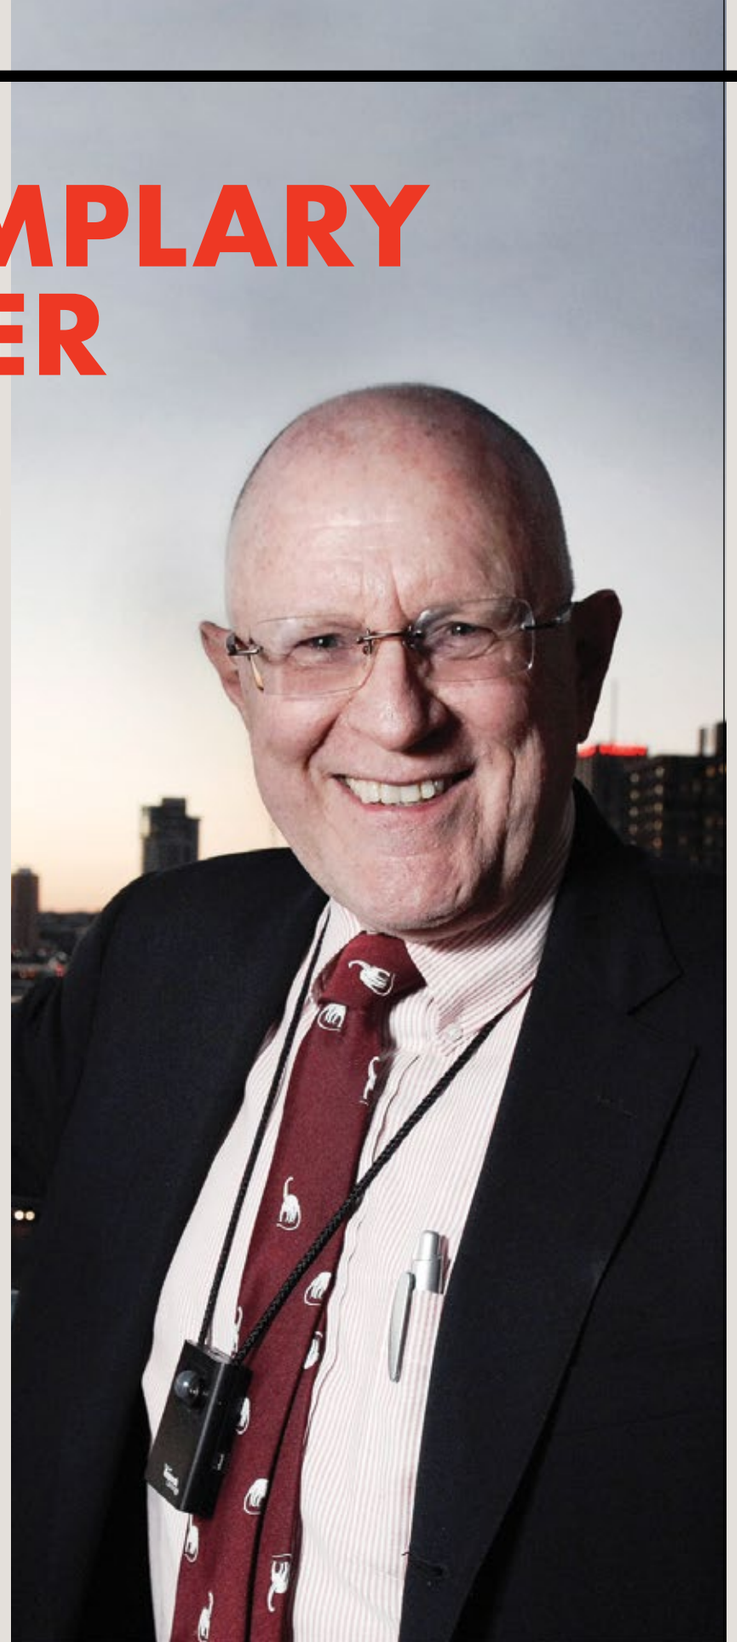

Think about the smartphone in your hand and compare it to the first iPhones that were launched in 2007. In the decade that has passed, the speed of your phone's central processing unit (CPU) has quadrupled, its storage has increased 16-fold, and it now boasts 20-fold more memory—all without costing very much more.

One man predicted this would happen. In 1972, Professor Gordon Bell formulated what is now known as Bell's Law, describing the formation, evolution and death of different types of computing systems. The Law predicted that computer classes would evolve along one of three paths depending on their price: an established class with a constant price but continually improving performance; a supercomputer class which would get progressively more expensive in the race to be the fastest; and finally, a cheaper 'minimal' class that would open up new markets.

More than 40 years after it was first proposed, Bell's Law continues to hold, with smartphones as an example of an established class that conforms to the Law.

"I believe it will persist for many decades to come, with different materials and transducers," Bell told *Supercomputing Asia*.

HANDS-ON EXPERIENCE AND HISTORICAL EVOLUTION

Computers were still in their infancy when Bell graduated with a degree in electrical engineering from MIT in 1957, but what he saw had him hooked. While in Australia on a Fulbright scholarship, he had the opportunity to work on English Electric DEUCE, a machine that Alan Turing had been involved with during his time at the National Physical Laboratory.

Back at MIT for his PhD studies, Bell worked on TX-0, one of the first computers to use transistors instead of the prevailing vacuum tube technology. It was then that he met fellow engineers Ken Olsen and Harlan Anderson, who persuaded Bell to join them in their newly formed company to work on a new prototype that they called Program Data Processor (PDP) minicomputers.

"What fascinated me was the incredible versatility to do everything including calculating, storing, communicating, controlling, etc.," he enthused. As the lead architect of the system, Bell helped the PDP series become one of the most popular minicomputers of all time, with the PDP-11 selling over 600,000 units.

That achievement was surpassed during his tenure as vice-president of engineering at Digital Equipment Corporation (DEC), where he spearheaded the Virtual Address Extension (VAX) line of microcomputers that catapulted DEC to the position of second largest company in the industry.

In between developing those two blockbusters, Bell taught computer science at Carnegie Mellon University from 1966 to 1972. While working on a book with his mentor Allen Newell, he became interested in the classification and evolution of different classes of computers. Coupled with his hands-on experience in developing successful microcomputers, Bell's academic interest in the historical evolution of the first computing systems culminated in the formulation of Bell's Law.

OF POLICY AND PRIZES

Bell's career took another turn in 1983, when a heart attack prompted his resignation from DEC. Far from slowing down, however, he turned his attention to computing at the national and international level, becoming the founding assistant director of the National Science Foundation's directorate for Computer & Information Science & Engineering (CISE) in 1986 and launching the ACM Gordon Bell Prize the following year.

"I think we will see IoT-based low power wireless devices integrated to form single chips for connecting to everything in the future."

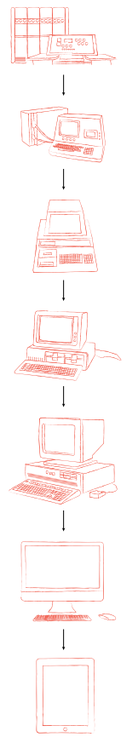

"I funded the prize to reward and acknowledge those people who program these highly parallel computers," Bell added. "The first prize really got the community interested in exploiting parallelism."

More than the US$10,000 cash reward, the Gordon Bell Prize is highly coveted for the prestige that it confers on the recipient, being akin to the Nobel Prize of supercomputing. The inaugural prize in 1988 went to Dr. John Gustafson, now with the Agency for Science, Technology and Research, Singapore.

Bell himself is no stranger to accolades, having won the National Medal of Technology in 1991 and the inaugural IEEE John von Neumann Medal in 1992, among others.

THE ACT OF MEMORY

Now in the third act of his career, the 82-year-old Bell continues to investigate the possibilities afforded by computers. In particular, he is interested in the possibility of using computers as an aid to human memory.

"I got the idea when [artificial intelligence expert] Raj Reddy wanted to scan the books I had written as his project to preserve the 20th century. I had boxes of all the papers, books, CDs, photos and videotapes that we were trying to document," Bell shared.

That project evolved into MyLifeBits, a real-time lifelogging experiment that attempted to capture every bit of information that Bell generated between 1998 and 2007. Microsoft, where Bell worked full time from 1995 till his retirement in 2012, has since developed software that annotates all that data, making it easy to sort and retrieve.

Though lifelogging is no longer popular, wearable technology and the Internet of Things (IoT) could spark a revival. "I think we will see IoT-based low power wireless devices integrated to form single chips for connecting to everything in the future. The next decade will probably be more deployment followed by all kinds of efforts to understand and exploit all the data," Bell said.

But at the end of the day, Bell does not see himself as a futurist. Of all the different hats he has worn over the course of his six decade long career, one remains closest to his heart.

"The early days when I actually did designs were perhaps the fondest memories," he said. "I would like to be remembered as an engineer who worked, by example, to support the creation of more engineers and engineering thinking." S

GREEN IS THE NEW BLACK

Satoshi Matsuoka on energy-efficient supercomputers

Supercomputers are power guzzlers, but researchers like Professor Satoshi Matsuoka are leading the charge for greener, more energy-efficient machines.

By **Rebecca Tan**

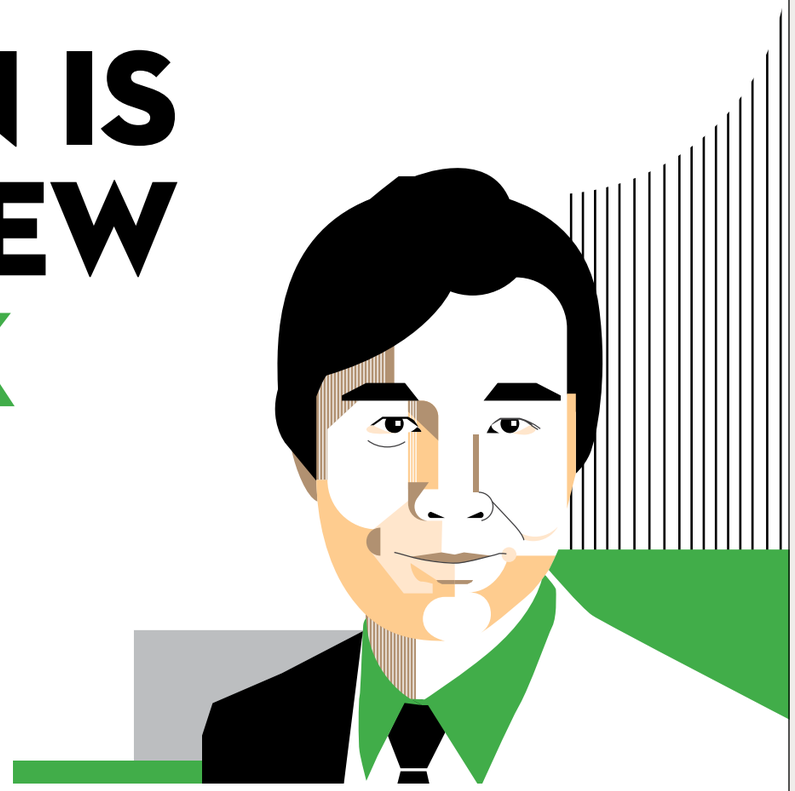

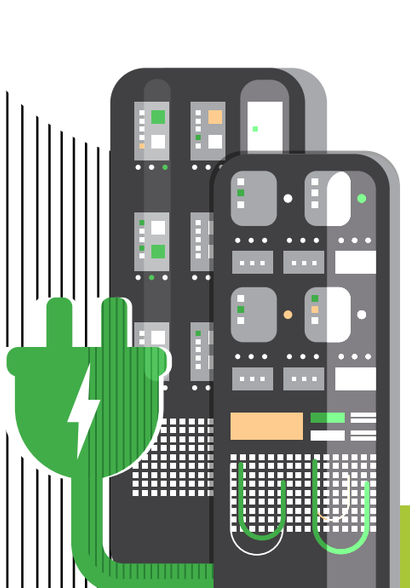

If supercomputers were cars, Tsubame-KFC at the Tokyo Institute of Technology would be a Tesla Model S—powerful but green. When it was unveiled in 2013, Tsubame-KFC promptly swept up first place on both the Green500 and Green Graph500 lists, convincingly demonstrating that it is possible to be mean, lean and clean. But just how its team achieved this feat might come as a surprise: they dipped their entire supercomputer in a vat of oil.

In this interview, Professor Satoshi Matsuoka, the key architect of the Tsubame projects, shares with *Supercomputing Asia* their seemingly unconventional approach to going green and the role that energy-efficient supercompters will play in the race to exascale.

Why is green supercomputing important?

Satoshi Matsuoka: The power budget has now become one of the most limiting constraints in high performance computing (HPC) and other large-scale information technology (IT) infrastructures.

In these days of massive parallel and distributed computing, machine performance is dictated by the size of the machine, which is proportional to the demand for energy.

Given that large machines are hitting megawatt-scale and beyond—equivalent to the electricity consumption of hundreds to thousands of households—increasing performance would be very difficult unless we reduce power consumption.

What makes Tsubame-KFC so energy efficient?

Tsubame means barn swallow in Japanese but is also an acronym which stands for Tokyo-Tech Supercomputer and Ubiquitously Accessible Mass-storage Environment.

KFC stands for Kepler Fluid Cooling.

SM: KFC's efficiency comes from various technological elements. One is the aggressive adoption of many-core technologies in the form of densely packed graphics processing units (GPUs).

Another is intricate power measurements and proactive control of the voltage and frequency of the GPUs to achieve maximum efficiency. One other major element is oil immersive cooling, by which we significantly reduce the power required for cooling, due to superior thermal conductivity of fluid versus air. [Oil has a specific heat capacity 1,200 times higher than air.] Oil immersive cooling further eliminates the need to power expensive fans, both in the chassis and the rack.

Also, although the oil may be at 40 degrees Celsius or more, all the components—especially the GPUs and central processing units (CPUs)—run at much cooler temperatures than air-cooled systems, again due to the massive difference in thermal conductivity. This reduces the leakage current.

Finally, since the temperature of oil can be higher than that of ambient air, it can be cooled in ambient air, rather than by using power-expensive heat pumps. Overall, both the machine itself and the cooling system are very efficient.

In which areas does Japan lead the supercomputing world? What are its comparative weaknesses?

SM: Japan's strength is that it has been involved with supercomputers for a long time and has experience with developing talent. We also have two vendors that are developing their own CPUs for HPC, namely Fujitsu and NEC. This gives us the capability to build the entire system from ground up, for both hardware and software. Also, there are many home-grown applications.

The weakness is that we have relied on home-grown technologies for too long. The IT system development is much more horizontal, assimilating technologies from everywhere, as demonstrated in other IT sectors such as smartphones. Also, the CPUs made by Fujitsu and NEC are not sold separately but only as a complete system, unlike those made by Intel or NVIDIA.

"Increasing performance would be very difficult unless we reduce power consumption."

Will Japan be the first to develop an exascale supercomputer? What are the key challenges standing in its way?

SM: I cannot comment on this publicly due to my non-disclosure agreement with the [Japanese] government regarding the Flagship 2020 "Post-K" project. What I can say is that in any exascale development the principal problem will be power efficiency, in addition to many other obstacles such as reliability, scaling, programming, I/O (input/output), etc.

Why has Japan made exascale computing a priority?

SM: There are many applications that will benefit from exascale. For Post-K, there are nine strategic application areas and four emerging areas. They cover everything from medical and pharma, environmental and seismic, climate and weather, manufacturing, advanced materials, brain simulations, combustion, as well as fundamental sciences such as the simulations of dark matter.

What technical advances in supercomputing do you anticipate in the next ten years?

SM: In the short term, I believe there will be a significant push for artificial intelligence (AI) and Big Data to become one of the dominant themes in HPC, steering the improvement of the entire ecosystem.

In the longer run we are approaching the end of the so-called Moore's law, around 2025 at the earliest. At that point, semiconductor lithography will no longer shrink, making transistor density and power constant over time, which means that FLOPS will cease to increase. This limit will be the most challenging problem facing all IT, not just HPC.

My belief is that the alternative strategy to increase performance beyond the end of Moore's law will be the relative increase in the performance of data-related parameters such as memory capacity and bandwidth, as well as the interconnect bandwidth; in other words, improving performance as measured by BYTES.

Along with the customization of processing to save transistors depending on the type of data, I see a transition of computing from FLOPS to BYTES. Of course, this entails a complete change in the use of devices, hardware architectures, software stack, programming, as well as algorithms. It will be a ten-year or longer challenge to devise a new paradigm for computing systems, a challenge that has to start now in basic research.

Business Bytes

BAIDU TURNS TO FIELD PROGRAMMABLE GATE ARRAYS

Chinese web services giant Baidu has chosen American chip maker Xilinx to supply field programmable gate arrays (FPGAs) to its data centers. The FPGAs will be deployed in Baidu's data centers for machine learning applications, such as in image and speech recognition.

The announcement follows hot on the heels of Microsoft's move to FPGAs for its data centers under the Project Catapult initiative, and marks a growing interest in FPGAs among the world's largest data centers.

Xilinx is known for inventing the first commercially-viable FGPAs, which are integrated circuits that can be configured by customers after manufacturing. Although harder to program than standardized, mass produced central processing units (CPUs), FPGAs are much faster than CPUs and offer greater flexibility than purpose-built application-specific integrated circuits (ASIC).

"Acceleration is essential to keep up with the rapidly increasing data center workloads that support our growth," Mr. Liu Yang, executive director at Baidu, as quoted by *DatacenterDynamics*. Director of Baidu's Autonomous Driving Unit Mr. Bao Junwei added that Xilinx FPGAs could potentially aid in the design of autonomous vehicles.

KEPPEL KEEPS THINGS SHIPSHAPE WITH COMPUTATIONAL FLUID DYNAMICS

Keppel Offshore & Marine and the National Supercomputing Centre Singapore (NSCC) have signed a memorandum of understanding to collaborate in the areas of supercomputing and high performance computing.

The March 2016 partnership agreement also includes collaborative opportunities in long range InfiniBand connectivity, networking and scientific software applications, green technology for data centers and related areas.

Headquartered in Singapore and one of the largest offshore and marine groups in the world, Keppel designs, builds and repairs offshore rigs, as well as a diverse range of vessels.

Keppel's work would not be possible without a deep understanding of fluid dynamics, the science of fluids in motion, said Mr. Aziz Merchant, executive director of Keppel Offshore & Marine Technology Centre, in an interview with *Supercomputing Asia*.

"Keppel optimizes the designs of rigs and vessels by tapping on one of the most advanced numerical analysis tools, computational fluid dynamics, which is able to numerically capture most environmental loads imposing on a vessel, as well as the vessel's response," Merchant said.

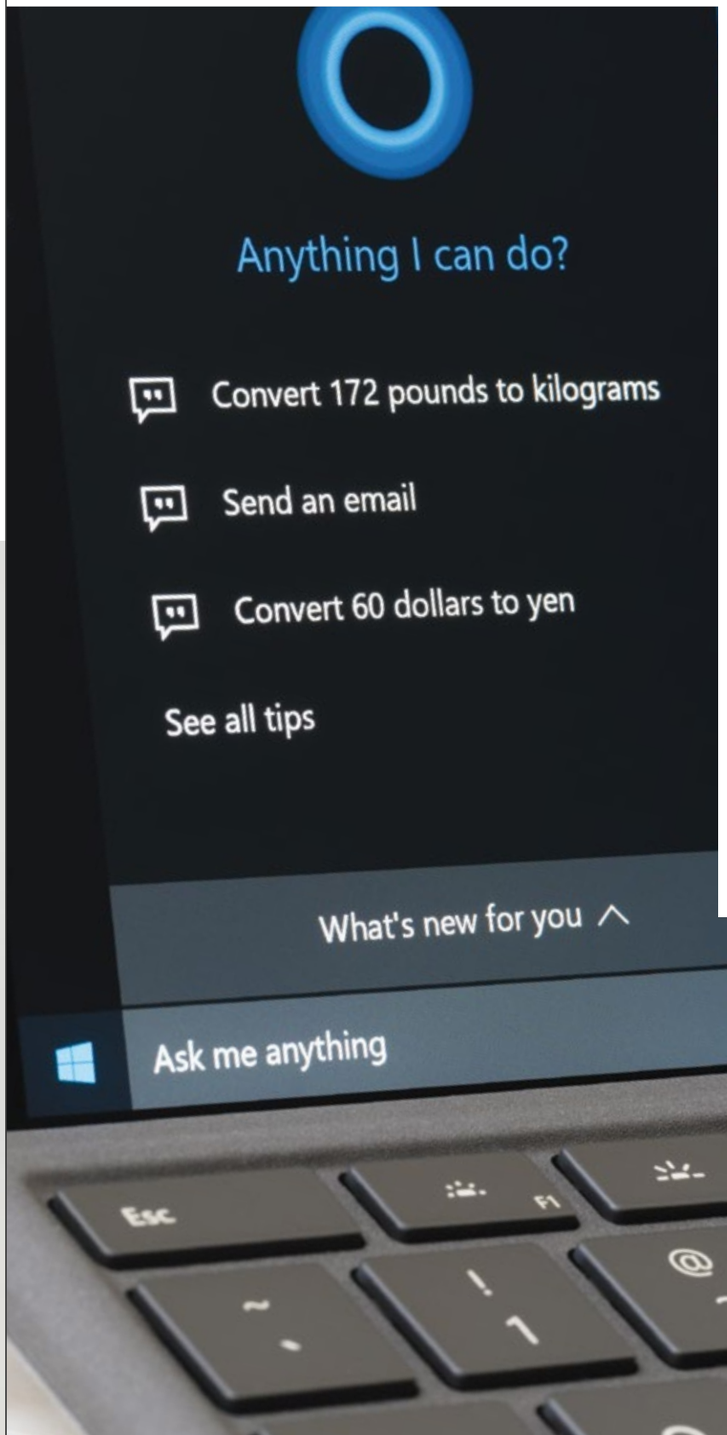

MICROSOFT SPEECH RECOGNITION REACHES HUMAN PARITY

In a culmination of decades of research, a team of researchers at Microsoft has created speech recognition technology that can recognize words in a conversation as well as any human.

In their paper, the team from Microsoft Artificial Intelligence and Research reported that their speech recognition system makes as many or fewer errors than a professional transcriptionist. What this means is that the rate at which the system misheard words, with a word error rate of just 5.9 percent, is the same as you would expect from a person listening to the same conversation.

Dr. Geoffrey Zweig, who manages the Speech & Dialog research group at Microsoft, attributed the accomplishment to the use of Microsoft's Computational Network Toolkit, a system for deep learning, in all aspects of the system. In deep learning, large amounts of data are used to teach computers to recognize patterns from input such as images or sounds.

"The next frontier is to move from recognition to understanding," Zweig said.

Super Snapshot

THE RACE TO BUILD AN EXASCALE COMPUTER

As countries around the world race to build the first exascale supercomputer, here are some numbers to give you a sense of scale.

EXA 10^{18}

[unknown]

1 petaFLOPS
ASPIRE 1 (2016)

PETA 10^{15}

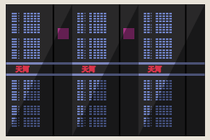

34 petaFLOPS
Tianhe-2 (2014)

93 petaFLOPS
Sunway TaihuLight (2016)

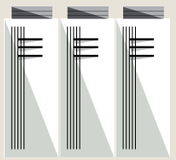

1 teraFLOPS
ASCI Red (1996)

TERA 10^{12}

1,000 teraFLOPS
IBM Roadrunner (2008)

0.1 gigaFLOPS
Cray-1 (1976)

GIGA 10^{9}

~70 gigaFLOPS
iPhone 6 (2014)

3 megaFLOPS
CDC6600 (1964)

MEGA 10^{6}

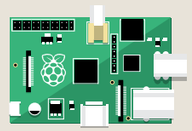

42 megaFLOPS
Raspberry Pi (2012)

More HPC Power. Lesser Hassle.
Speed up your research and development with HPC.

You are the key to solving complex scientific and technical problems to create a better future. We provide you with the computing power to make the right decision.

Fujitsu High Performance Computing technologies combine performance with superior energy efficiency to help you increase productivity and boost performance. From hardware to software, we provide high levels of reliability and operability, flexibly configured to meet every need.

For more information, contact info@sg.fujitsu.com.

shaping tomorrow with you

FUJITSU