NSCC
ASPIRE

ISSUE 1 I MAY 2017

Introducing NSCC ASPIRE

Welcome to the inaugural issue of NSCC ASPIRE, a quarterly e-publication that keeps you updated on the latest developments in high performance computing (HPC) in Singapore as well as exciting use cases in education, industry and research.

Titled after ASPIRE 1, Singapore's first petascale supercomputer, the e-newsletter is part of efforts by National Supercomputing Centre (NSCC) Singapore to bring the benefits of HPC to the wider community. We are excited about the opportunities that lie ahead and we would like to invite you to join us as we embark on this journey towards the democratisation of HPC.

– NSCC Communications Team

HPC luminaries light up Supercomputing Frontiers 2017

Mr Peter Ho

Prof Gordon Bell

Prof Thom H. Dunning Jr

Prof Fu Haohuan

Prof Alessandro Curioni

PROF GORDON BELL of the Gordon Bell Prize; Prof Fu Haohuan of National Supercomputing Centre in Wuxi (NSCC-Wuxi) China and Prof Alessandro Curioni from the IBM Research Laboratory in Switzerland, both of whom share the distinction of being Gordon Bell Prize winners; and Prof Thom H. Dunning Jr, who led the US National Centre for Supercomputing Applications – these were some of the big names that lit up Supercomputing Frontiers Singapore 2017.

In their conference keynotes, the luminaries from the world of high performance computing (HPC) shared their views on the challenges and opportunities in HPC and gave the audience a rare glimpse into the development of some of the most powerful computing machines in the world.

Welcoming these insights, Mr Peter Ho, Steering Committee Chairman of the National Supercomputing Centre Singapore (NSCC), said Singapore was only at the start of developing and implementing an HPC strategic roadmap. "For the NSCC to realise the vision of catalysing supercomputing in Singapore, it still will have to bootstrap on the experience and knowledge of others."

ASPIRE 1

It was only in December last year that Singapore entered the petascale league with the official launch of the national Advanced Supercomputer for Petascale Innovation, Research and Enterprise or ASPIRE 1, said Mr Ho.

The first petascale supercomputer in Southeast Asia, ASPIRE 1 will underpin the national research infrastructure of Singapore's Research,

Prof Gordon Bell recounted the history of his eponymous supercomputing prize in a presentation on "Motivating and Measuring HPC Progress".

Innovation and Enterprise masterplan or RIE 2020.

Run by NSCC, it sits on top of the national InfiniBand fabric that extends to the campuses of NSCC stakeholders including the Agency for Science (A*STAR), Technology and Research; National University of Singapore (NUS); Nanyang Technological University (NTU); and the Singapore University of Technology and Design (SUTD). It is also connected to an international network for the deployment of InfiniBand global range. This gives ASPIRE 1 the potential to connect with any supercomputer centre in the world for collaboration and cooperation.

Supercomputing centres in Australia, Catalonia, Japan, Europe and North America are already piloting links to ASPIRE 1 either via the NSCC Science, Technology and Research Network (Star-N) connectivity project or the InfiniCortex project, a worldwide InfiniBand fabric that brings together several supercomputing facilities spanning Asia, Australia, Europe and North America.

Mr Ho noted that at its launch last year, ASPIRE 1 was already operating close to maximum capacity. "That was a very encouraging start," he said. "To the NSCC, it is an indication of future demand and the strong promise of supercomputing in Singapore."

At the Supercomputing Frontiers event, NSCC announced the soft launch of payment plans for the use of ASPIRE 1. These ranged from $500 to $5,000 per month depending on requirements such as the number of user accounts, active jobs per tier and maximum CPU cores per job.

Inaugural HPC Awards

Also making its debut at the event was NSCC's HPC Awards which recognise outstanding individuals who have contributed to Singapore's HPC landscape as well as research and commercial efforts that tap on ASPIRE 1's computational power to drive innovation, raise productivity and improve lives.

One of the winning teams comprised Dr Anjan Soumyanarayanan and Dr Anibal Gonzalez from A*STAR's Data Storage Institute (DSI), who won the NSCC Outstanding HPC Scientific Award for their study on particles known as skyrmions that could pave the way for applications in memory devices of the future.

"Materials and device-based research benefit greatly from service-based supercomputing access," said Dr Anjan, Scientist from DSI. "Service-based HPC enables the use of well-developed software packages for rapid simulation of material properties and device behaviour with realistic configurations. This leads to an accelerated, yet nuanced understanding of performance limitations and avenues for improvement towards developing real-world technologies."

The research team used the NSCC resources to simulate the impact of magnetic properties of thin film materials on the behaviour of particles known as skyrmions. Through this effort, the team demonstrated the presence of skyrmions with tunable properties in multilayer films compatible with industrial fabrication processes. "Such a tunable skyrmion platform can pave the way for their use in future memory devices," said Dr Anjan.

Demonstrating the wide range of applications that are possible with HPC, the award submissions also included Keppel Offshore & Marine Technology Centre's use of NSCC's HPC resources to optimise the designs of rigs and vessels in the product development process; and Tan Tock Seng Hospital's use of HPC in combination with whole-genome sequencing to determine the transmission routes of bacteria and provide useful information for on-the-ground infection control.

Singapore Supercomputing Frontiers 2017 was held in Biopolis on 13-16 March. The conference and exhibition attracted some 450 participants from 12 countries. ■

Winners of the inaugural NSCC HPC Awards

Supercomputing Frontiers Singapore Distinguished Service Award
Mr Lim Chuan Poh, Chairman, A*STAR

Supercomputing Frontiers Singapore Visionary Award
Dr Raj Thampuran, Managing Director, A*STAR

NSCC Outstanding HPC Scientific Award
Data Storage Institute, A*STAR

NSCC Outstanding HPC Industry Application Award
Keppel Offshore and Marine Technology Centre
Honourable Mention: Global Gene Corp Pte Ltd

NSCC Outstanding HPC Innovation Award
Institute of High Performance Computing, A*STAR
Honourable Mention: Singapore Centre for Environmental Life Sciences Engineering, Nanyang Technological University

Best Integrated Vertical Domain (HPC in Offshore & Marine)
Special Honourable Mention:
- Technology Centre for Offshore and Marine Singapore
- Department of Fluid Dynamics, Institute of High Performance Computing, A*STAR
- Keppel Offshore and Marine Technology Centre


Singapore Distinguished Service Award – Mr Lim Chuan Poh.


Singapore Visionary Award – Dr Raj Thampuran.


NSCC Outstanding HPC Scientific Award
– Data Storage Institute.


NSCC Outstanding HPC Industry Application Award
– Keppel Offshore and Marine Technology Centre.

Solving supercomputing's software challenge

Porting scientific models over to new heterogeneous supercomputers remains a daunting task.

EVEN THE MOST powerful supercomputer in the world could not overcome an aversion to writing climate modelling programs.

This was what Professor Fu Haohuan from the National Supercomputing Centre in Wuxi (NSCC-Wuxi) found out when working with students on a climate modelling project using the 98-petaFLOP, 10,649,600-core Sunway TianhuLight.

Prof Fu, who is also from the Department of Earth System Science at Tsinghua University, was a member of the 2016 Gordon Bell Prize-winning team that succeeded in simulating the vast number of variables inherent in a developing weather system, to predict and simulate the weather and related atmospheric events.

At the start of the project in 2015, there were 21 students on the team.

By the time it ended, only two or three remained. And when Prof Fu tried to entice a promising young man to stay on with the prospect of working on the most powerful supercomputer in the world, the student was keen but had one condition – that he did not have to deal with climate code.

Prof Fu was not surprised. "Many of the students are trained in computer science and they do not have a good understanding of the physics and numbers behind climate code. The code is also difficult to read," he said.

"If we are all facing the challenge of porting different sets of scientific code onto new supercomputers, we probably have to start with the education programme in the university, to try to make it more interdisciplinary."
— Prof Fu Haohuan

Multi-disciplinary approach

Recounting this experience at the Supercomputing Frontiers 2017 conference, Prof Fu said the lack of people/talent was the biggest challenge facing supercomputing projects. "There is a misfit between the in-place design philosophy and the new architecture, and a lack of people with interdisciplinary knowledge and experience."

Presenting on the topic "Designing and Tuning Scientific Applications at the Scale of 10 Million Cores", Prof Fu noted that many of the prevailing scientific models consist of millions of lines of legacy code written for a multi-core rather than a many-core architecture, and they scale poorly.

The code would have been written by different teams and have different code patterns, and it would be difficult to force them into different architectures, especially if they had been designed for a multi-core architecture, he said.

Prof Fu observed that in China, there were many very large heterogeneous systems, but many existing supercomputing applications such as the climate models were not ready to take advantage of these high performance capabilities.

To address this and enable climate models to benefit from new supercomputer systems, NSCC and Wuxi are collaborating to build new, highly scalable frameworks that can efficiently utilise the new architecture with the many-core processors. The effort involves a team of people from different disciplines such as computational mathematics, computer science and geo-computing.

In Prof Fu's view, the multidisciplinary ethos is important, and

"When real performance deviates from peak performance, it is a software issue. Substantial advances in performance are only possible if there are new algorithms or approaches to overcome these system limitations."
— Prof Thom H. Dunning Jr

it has to be cultivated early in a researcher's career. "If we are all facing the challenge of porting different sets of scientific code onto new supercomputers, we probably have to start with the education programme in the university, to try to make it more interdisciplinary," he said.

This would be one way of closing the gap between hardware and software in order to take full advantage of the capabilities of next-gen supercomputers.

Closing the gap

Speaking on "The Opportunities and Challenges of Exascale Computing", Prof Thom H. Dunning Jr from the Northwest Institute for Advanced Computing in the United States pointed out that real world performance depends on applications taking advantage of hardware innovations. "We have to worry about the software," he said.

Prof Dunning highlighted the fact that over the years, despite an increase in performance capabilities in high performance computing, efficiency has declined from 30-40 per cent on vector supercomputers of the 1990s to as little as 5 to 10 per cent on parallel computers today.

"When real performance deviates from peak performance, it is a software issue," he said. "Substantial advances in performance are only possible if there are new algorithms or approaches to overcome these system limitations."

The software challenge in HPC was recognised by Prof Gordon Bell in the 1980s, when the Gordon Bell Prize was introduced to recognise outstanding achievements in high performance computing applications.

Recounting the history of HPC in a presentation on "Motivating and Measuring HPC Progress", Prof Gordon Bell said the intent of the prize was to acknowledge the "incredible difficulty" involved in writing programmes that can take advantage of parallelism.

Three decades on, the programming aspect of supercomputing continues to be a challenge but it is still being overlooked in the hardware-chauvinist world of HPC. As Prof Dunning observed, "There is a lot of focus on hardware and not nearly enough on the software side."

There has to be a more integrated approach to advancing computing, he said, "You need to also worry about the node OS, runtime, what the system software looks like. You need to consider if new programming models or new maths libraries will be required. When configuring a supercomputer, you also have to pay attention to the scientific problem that you are trying to solve."

He cited the example of the quest by the United States Department of Energy to develop an exascale computer together with industry.

"They are not looking to build one computer for one application but an exascale computer designed such that a whole wide range of applications can run on it, whether it is chemistry, climate and geophysics or astrophysics or a whole range of other scientific applications. The whole goal of the project is to, as much as possible, integrate all these activities together."

Weighing the trade-offs

The important question to ask, therefore, is, "What do you want?"

Taking up on this point, Mr Joseph Curley of Intel USA noted that in today's marketplace, there are different types of computers ranging from general purpose machines to specialised machines, and that "applications built for one may not be built for the other".

Sharing his views on "Applying Petascale Learning to Exascale", Mr Curley pointed out that optimisation choices for one application often come at a cost for general applications.

The programming models are different. For example, it could be general purpose code running procedural language on a general purpose machine, versus building an application to expose the capabilities of the parallel machine.

A mismatch could lead to problems of performance and productivity, such as underutilised CPUs. This is an "opportunity gap", said Mr Curley. "Fully utilising existing hardware for both performance and power efficiency is an imperative, and that is where the focus should be on."

Echoing the point made by Prof Dunning, he said, "Great attention is being paid to hardware on the path to exascale. We need as great a focus on the applications." ■

Transforming R&D with HPC and AI

The convergence of supercomputing and big data could disrupt traditional R&D and reduce the research process from years to weeks.

Prof Alessandro Curioni: The basis of research could shift from opportunistic discovery by humans to comprehensive discovery by cognitive computing.

ARTIFICIAL INTELLIGENCE (AI) and big data will amplify supercomputing's role in R&D and provide a "disruptive opportunity" that could reduce the research process from years to weeks.

This was one of the possibilities discussed at Supercomputing Frontiers 2017, where luminaries from the world of supercomputing shared the stage with industry experts to discuss the challenges and opportunities of AI in the context of high performance computing (HPC).

In his keynote on "Cognitive Discovery: The Next Frontier of R&D", Prof Alessandro Curioni from the IBM Research Laboratory in Switzerland noted that supercomputing had made a great contribution in establishing simulation as the third pillar of R&D, the other two being theory and experimentation.

"We have been able to push the ease of use of simulation, to increase simulation accuracy and fidelity, to have better scalability on parallel machines and to get more computational power and enhance throughput – all these thanks to supercomputing."

However, traditional discovery has its limits. "We need a new data-driven, holistic approach," he said.

Traditional approaching

The traditional approach to R&D has been to do things in a linear way, from opportunistic discovery to simulation and experimentation. As Prof Curioni described it, it goes something like this: "I know something, I do an experiment, I run a set of simulations, I take one that I need, I do another experiment, and that goes. on. It is a very lengthy process."

However, the data explosion that is taking place today could transform this. The volume of data in the world is predicted to reach 165 zettabytes by 2025. In R&D, the amount of knowledge that is generated from research is also growing at an exponential rate. Citing an example from the field of materials science, Prof Curioni noted that the number of papers published in this field has grown from a few a week to 50 or more in a single day.

"This is a huge amount of information or knowledge that is difficult for us to make sense of, and without it, we are limiting ourselves to our own knowledge," he said, recounting how, as a researcher, he could read maybe five papers at the very most in a good week. "So you have limited knowledge and based on this limited knowledge, you build an empirical simulation model using your experience only. The process takes months, and then lab tests take months to years. To go from the initial thinking to the results you need years."

The rise of cognitive systems

One approach to resolve this has been the creation of cognitive systems powered by high per-formance computing to augment human cognition. These systems are able to take the tremendous amount of knowledge that is produced in different fields and integrate it in real time, learn from the data, and contribute the insights back into the pool of human knowledge. “This connection, putting things in context, magnify by order of magnitude the value of the data,” said Prof Curioni.

“You start to ingest all the knowledge using cognitive machines and machine learning, and you can do it in days. Once you have it, you create some sort of knowledge graph that you can query. You can do inference, and this will allow you to understand which direction you need to go, and the development process can go from years to weeks.”

With this, researchers will be able to use AI to target which simulations to run and the best way to execute them.

“The machine will be able to tell us what is the important knowledge within all that data that will help us solve our problem, and also which simulation we need to run to derive additional knowledge to what we already know. It will tell us not only which simulation we have to run, but also amongst the different simulation methods, which is best for that particular problem. This will reduce the number of simulations that you run and the number of experiments that you need to do.”

This shifts the research process from one based on opportunistic discovery by humans to one based on comprehensive discovery by cognitive computing. In this new world of R&D, theory is built on the global knowledge that has been accumulated, and simulation can go together with inference that is based on this knowledge, and “R&D goes from a linear process to something that is fully embedded”.

Mr Cheng Jang Thye: The convergence of analytics and HPC will enable different types of simulation to be carried out.

The IoT opportunity

Beyond the realm of research, another area where supercomputing looks set to have an impact is in the Internet of Things (IoT).

Presenting on the topic “Leveraging on HPC for IOT and Big Data Analytics”, Mr Cheng Jang Thye, Chief Solutions Architect of Fujitsu Singapore said the cloud will allow HPC technologies to be shared by more people. At the same time with IOT, the data will be coming in live, streaming in from many sources.

As these developments converge, he sees HPC algorithms being used in areas such as edge analytics – “boxes on roads and lamps posts with CPU cores” – to analyse data that is coming in, in real time.

“The scope of HPC is much larger than what most people think,” he said. The convergence of analytics and HPC will enable larger models to be built and different types of simulation to be carried out using different algorithms.

“These kinds of computation can be used for different kinds of problems that we see in the field,” he said, citing examples such as fraud detection, anomaly detection, and helping marketers to understand customer patterns.

HPC is also fundamental to deep learning, which is reaching “viable accuracy”, said Mr Cheng.

He cited AlphGo, an artificial intelligence program developed to play the board game Go, as a proven result of deep learning using a lot of computation in a cluster model and algorithms designed to meet specific user cases. “It is the algorithm that creates the intelligence, and there are a lot of opportunities to use these kinds of capabilities to accelerate deep learning applications,” he said.

Scaling up and scaling out

The use of HPC in AI requires scalability on different fronts. Presenting on “Supercomputing for AI”, Mr Marc Hamilton, Vice President of Solutions Architecture and Engineering at NVIDIA noted that the “deep learning” aspect of AI is an “exaflop computing problem”.

“It is about bringing in millions of images, running them through a deep learning neural network and developing the model,” he said.

Once the model is trained, it can be pointed to an inference application which could be a web application or a mobile application, or it could be something like a self-driving car. And that is a “scale out” model involving potentially millions of cars.

“But training for that is a scale-up,” said Mr Hamilton. “You can’t build a self-driving car without a scale-up deep learning system.”

In this respect, there is a lot that the HPC community can teach the people in the AI world, he said. “Because increasingly, AI is supercomputing”.

The "scaling out" of HPC

Access to NSCC's ASPIRE 1 supercomputer enables students from National Junior College to investigate various two-dimensional materials without having to deal with computational constraints.

The NJC students with NSCC Steering Committee Chairman Mr Peter Ho and Director A/Prof Tan Tin Wee at Supercomputing Frontiers 2017.

THE TRADEMARK RED blazers of National Junior College (NJC) stood out at the Supercomputing Frontiers exhibition in more ways than one. Amidst industry partners presenting high performance computing (HPC) solutions and researchers explaining various scientific use cases, the students were holding their own, attracting the attention of no less than Prof Gordon Bell of supercomputing's Gordon Bell Prize.

Over the past year, a tie-up with the Singapore National Supercomputing Centre (NSCC) had given the NJC students access to NSCC's ASPIRE 1 petaflop supercomputer – a machine which is capable of processing one thousand million million (1015) floating-point operations per second. This allowed Mr Harman Johll, Head of Department/Research at the junior college, to mentor students who were interested in computational materials science and to investigate almost any material that they might be interested in, without having to deal with computational constraints.

"Parameters that I had previously 'worried about', such as the energy and energy density cutoffs and the number of k-points in the Brillouin Zone, were no longer of any concern," he said. "Simply calibrate and use what the results tell you without compromise on system size. This allowed us to do fantastic work, fast – critical attributes to research that is publishable and that would add knowledge to the scientific community."

The capacity to explain atomic or molecular-scale processes and even to design materials with specific properties depends on developing strong theoretical models and having sufficiently powerful machines to apply those models, said Mr Johll.

With the computing power available to them, different teams of students were able to investigate various two-dimensional materials such as graphene, borophene and silicone – substances that have a thickness of a few nanometres or less – and how these could be functionalised for particular applications.

"The students were absolutely excited when I told them we would have access to NSCC (or HPC)," said Mr Johll. "They had already learnt the basics, for example, about Linux and the use of the density functional theory software package. All they had to do was learn how to use system scripts, something that took little over 10 minutes to teach."

The projects allowed the students to develop their sense of modelling, expand their knowledge of the materials world and methods available to understand systems at the atomic or molecular scale, and "perhaps most importantly, excite them in doing science", said Mr Johll.

For Prof Bell, who was a keynote speaker at Super-computing Frontiers event, the students' use ASPIRE 1 produced "probably the most exciting applications" that he saw at the event. It was a clear indication that HPC is "scaling out" in terms of the larger numbers of people using it, he said.

Brought to you by:

National Supercomputing Centre

Founding stakeholders:

Agency for Science, Technology and Research
Nanyang Technological University
NUS National University of Singapore
Singapore University of Technology and Design

National Supercomputing Centre (NSCC) Singapore | 1 Fusionopolis Way #17-01, Connexis South, Singapore 138632 | Tel: (+65) 6714 9450 | contact@nscc.sg