National Supercomputing Centre (NSCC) Singapore e-newsletter

NEWSBYTES

October 2021

In this Issue...

CORPORATE NEWS

NSCC and ITE equipping professionals with the basic knowledge of high performance computing

The first session of the Certificate of Competency (CoC) in Introduction to High Performance Computing (HPC) saw good turnout and feedback.

Held on 1 October 2021 at ITE College West and jointly organised by NSCC and ITE College West, the first session of the CoC in Introduction to HPC was attended by 15 participants comprising of ITE College West staff and lecturers. Course participants were co-trained by Dr Lee Thong Yan, HPC Application Analyst at NSCC, and Max Chua, Lecturer at ITE College West, on the basic building blocks of HPC and how to access HPC remotely from a virtual platform

Upcoming NSCC Webinar

Preparing Singapore for Climate Change

3 November 2021
3.00pm - 4.30pm (UTC+8)

Supercomputing for Climate-informed Urban Planning

by Dr Heiko Aydt,
Lead Investigator of Digital Urban Climate Twin R&D at Cooling Singapore,
Singapore-ETH Centre

Sea-Level Rise Emergency

by Prof Benjamin Horton,
Director, Earth Observatory of Singapore

LIVE panel discussion to follow

Moderated by Dr Dale Barker,
Director, Centre for Climate Research Singapore,
Meteorological Service Singapore

REGISTER HERE

to experience working on thousands of computing nodes to perform complex program tasks at high speed, which in turn will accelerate the building of deep learning AI applications.

Training accounts with computing resources were provided by NSCC. Upon completion of the course, participants were awarded a CoC in Introduction to HPC as well as a Certificate of Participation by NSCC Singapore.

NSCC and ITE will be holding another session of the CoC in Introduction to HPC on **20 December 2021** from **9am-5pm**.

At the end of this course, participants will acquire skills and knowledge on:
• Basic building blocks of a supercomputer
• Understanding PBS Job Scheduler
• Use-case & Accessing of HPC
• Environment Setup & File Transfer
• Resource Allocation & Job Submission
• Hands-on AI Project using HPC

"The CoC in High Performance Computing was an eye opening experience for many of the lecturers in ITE. We were able to experience first-hand what AI can do. I am looking forward to using HPC for more of our AI projects in the near future."

Pradeepkumar
Lecturer
Electronics, SEIT
ITE

Fees for this course can be paid using SkillsFuture Credits

Scan to register NOW!

tinyurl.com/58cbp9r9

For more information, head over to https://www.ite.edu.sg/courses/course-finder/course/coc-in-introduction-to-high-performance-computing-(hpc).

Back to main content list

Locked in arms

Dr Jernej Zidar
Senior HPC Application Analyst, NSCC Singapore

Not long ago the catchphrase in the gaming community was "But, can it run Crysis?". The question was referring to the first installment of the game Crysis released in 2007. While the image quality when roaming tropical islands was very realistic, the system requirements to run this game were so high that even very high end computers with overclocked processors (CPUs) and video cards struggled. Video (graphic) cards have come a long way since then, modern video cards contain a plethora of special hardware responsible for different functions that result in realistic imagery and smooth game play at high resolution.

This special hardware includes shaders, specialized units responsible for drawing the image that ultimately ends up on the computer screen. Nvidia later switched to using the term CUDA core to describe them. CUDA is Nvidia's proprietary API that enables parallel computing on graphic cards, open alternatives include OpenCL and SYCL from the Khronos group. CUDA and its alternatives enables applications to offload some or all compute tasks from the CPU to the GPU, these tasks can make use of the hundreds (on older hardware) and thousands CUDA high frequency cores to greatly accelerate all sorts of calculations. One issue is that these calculations must be parallelizable and it was something that slowed the initial adoption but once a critical mass of the algorithms and applications was able to make use of the untapped computational power the benefits of using GPUs became apparent. Two areas in particular greatly benefit from the use of GPUs: molecular dynamics and Artificial Intelligence/Machine Learning.

In molecular dynamics simulations Newton's equations of motion are solved for a biomolecular system (e.g. protein in an aqueous solution), this process is highly amenable to parallelization, including parallelization on GPUs. When some tasks are offloaded to GPUs, the performance of the simulation can improve by 2-3 times as compared to simulations running solely on CPUs. Molecular dynamics applications that can offload tasks to GPUs include Gromacs, NAMD and CHARMM. There are also efforts to port applications from the quantum chemistry domain so they can make use of GPUs as well. The partial success stories include Gaussian, NWchem, CP2K to name a few. Why partial? Quantum chemistry codes employ many different methods e.g. Hartree-Fock, Density Functional Theory etc. to solve or optimize a given group of atoms and not all methods can be readily parallelized to use GPUs.

A second area that greatly benefits from using GPUs are AI/ML workloads. The benefits of using GPUs are so great that these days CPUs are used only for training fairly small models, whereas the ever larger models are running either on GPUs or very specific hardware such as Google's Tensor Processing Units (TPUs). The large number of GPU cores can perform several computations at the same time, which is what happens in the training phase when several hundred or thousands variables have to be fitted to describe the model. Packages that can make use of GPUs include TensorFlow and PyTorch.

The two areas have more than just GPUs in common; in both areas the CPUs powering the compute nodes are still important but less and less so. Thus, it is possible to imagine a future where more and more tasks, including but not limited to molecular dynamics and AI/ML, are executed either on GPUs or some other specialized hardware, with the main CPU taking care of the scheduling and making sure all these specialized hardware units can communicate with each other without errors. In light of this, Nvidia's 2019 purchase of ARM, the designer of very efficient processors powering most if not all mobile phones, is a sensible one. If the vast majority of compute related tasks can be executed more efficiently and quickly on GPUs or similar specialized hardware, then the main CPU can be a much simpler one, something that is based on ARM, power efficient with a performance that is "good enough". The idea may appear far-fetched yet it has been embraced by makers of mobile devices, where each device - say a mobile phone - contains hardware units responsible for image processing, signal processing, machine learning and yet the main processor continues to be fairly simple.

A more plausible future is a multiplatform one where there is no solution that fits all problems. In this future there are perhaps GPU-powered compute nodes with a small ARM-based processors that are used solely for molecular dynamics or AI/ML, nodes that are powered by modern Intel or AMD processors for more generic workloads, nodes powered by ARM-based processors offering acceptable performance for tasks that don't require the ultimate performance. As it turns out, the computing landscape has been slowly moving towards this. ARM-based computers both personal and cloud-based are becoming more common and similar for other types of compute. The NSCC ASPIRE2A system, which users are anxiously waiting for, is buckling this trend as well with 768 nodes powered by AMD's Epyc processors and 82 nodes equipped with a total of 352 Nvidia A100 GPUs.

It is indeed a great time to be in high performance computing!

***Dr Jernej Zidar** is a Senior HPC Application Analyst at NSCC. High performance computing was instrumental for his diploma in biochemistry and the PhD in biomedicine. He's passionate about computers, whether on premise or in the cloud.*

Back to main content list

Training conventional surveillance systems to be smarter

Researchers at NTU are using HPC to design smart and autonomous cameras to aid in digital monitoring and crime prevention.

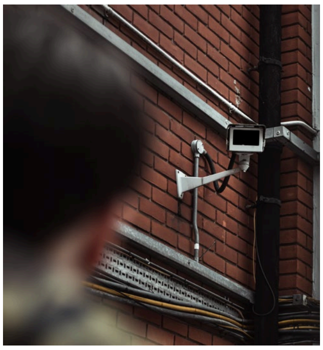

Video surveillance systems have been around for many decades and are installed all around the world. Such systems help monitor the targeted surroundings, alerts the authorities to incidents and act as a deterrent in helping prevent crime. These camera systems are also used by commercial owners to protect their assets.

With the recent success in applying deep learning in image and video processing, the roles of surveillance cameras are evolving. Law enforcement can now leverage intelligent and automated surveillance systems to perform reliable image and video analysis in real-time and with minimal human intervention. Examples of such applications include digitally analysing the movements of suspected criminals and terrorists, identifying and locating a missing person and detecting suspicious or abnormal pattens or behaviour to prevent potential crimes. However, older surveillance systems are not equipped with such intelligent capabilities and continue to rely on manual operators to be involved in the time consuming process of monitoring and analysis, which is both expensive and inefficient.

"We are leveraging NSCC's high-performance AI GPUs to efficiently train our large model and to explore many different design considerations to push the boundary. We are using the NSCC DGX system to handle the computationally intensive work of model training and hyperparameter tuning and as a result, are able to achieve state-of-the-art accuracy performance."

Tay Chiat Pin
Researcher
School of Electrical and Electronic Engineering
Nanyang Technological University

With advancements in machine visuals, and in recent years the rapid progress in the domain of computer visuals using deep learning, the demand for smart and autonomous cameras has risen sharply. These tasks can be largely managed by the person re-identification (re-ID) deep learning network with integrated facial and action recognition. To address this, a team of researchers at Nanyang Technological University (NTU) School of Electrical and Electronic Engineering are tapping onto NSCC's supercomputing resources to conduct research into person re-identification computer visual tasks to track human movement for safety and security purposes. The top two priorities of their work are to design a strong deep learning model that overcomes occlusion, lighting variation, and similar appearance problems and to work towards the quick deployment of the model in the field via unsupervised learning.

To find out more about how NSCC's HPC resources can help you, please contact e-news@nscc.sg.

Back to main content list

Accelerating research to curb the spread of metastatic cancers

Harnessing supercomputers to understand the molecular and cellular processes of metastatic cancers and to develop novel therapeutics targeting the disease.

Metastasis, which is the spread of cancer cells from its originating organ to distal organs, is the main cause of death in cancer patients. The World Health Organization (WHO) had projected colorectal and pancreatic cancers to be among the top ten contributors to mortality by 2030, with metastatic disease to be the main contributing factor to patient mortality in these cancers.

Metastatic diseases are usually refractory to standard-of-care therapeutics and not-amenable to surgical resection. Therefore, the five-year survival rate of patients with metastatic disease is dismal, with less than 15% of such patients expected to survive beyond five years. This points to an urgent need to identify therapeutic vulnerabilities in metastatic colorectal and pancreatic cancers.

A team of researchers from A*STAR's Genome Institute of Singapore is utilising NSCC's high performance computing resources to understand the underlying molecular and cellular processes of metastasis in order to develop novel therapeutics targeting metastatic disease. In recent years, single-cell sequencing technology has matured, which allowed for high resolution transcriptomic interrogation of the complex cellular milieu of metastatic cancers. Such sequencing efforts allowed researchers to identify the cellular heterogeneity inherent in cancers and understand how diverse cell types can interact cooperatively with each other to drive cancer progression despite treatment. Additionally, single-cell sequencing can also allow researchers to identify rare cancer cells that are inherently resistant to therapy and can initiate metastasis. By understanding the biology and source of metastatic relapse, new therapeutics can then be developed which either disrupt the synergistic inter-cellular interactions driving cancer progression or by specifically targeting metastasis-initiating cells.

"While single-cell sequencing is a very powerful tool for us to interrogate the cellular heterogeneity of cancers, it generates large amounts of sequencing data, which makes it impractical to perform the upstream data processing and analysis on our typical workstation. Being able to process the large amounts of sequencing data generated on NSCC's supercomputing resources has benefited us greatly. We were able to run memory intensive and/or multiple analysis in parallel, which sped up our workflow and the reliability of the cluster is much appreciated."

Loo Jia Min
Senior Research Fellow
Genome Institute of Singapore,
A*STAR

Niranjan Shirgaonkar
Bioinformatics Specialist
Genome Institute of Singapore,
A*STAR

To find out more about how NSCC's HPC resources can help you, please contact e-news@nscc.sg.

Back to main content list

THE LAST BYTE...

<SHARED CONTENT>

Shared articles and news from the HPC world.

NUS and Thales to develop quantum technologies

The National University of Singapore and Thales have joined hands to test quantum technologies for commercial applications in security and sensing.

Credit: Thales

The National University of Singapore (NUS) and Thales have teamed up to test quantum technologies for commercial applications. The partnership will see industry and academic experts from Thales and NUS's quantum engineering programme (QEP) develop capabilities to test and evaluate interdisciplinary quantum security technologies. They will also explore potential research collaboration opportunities in the fields of new materials and design for quantum sensing. Read more at Computer Weekly here.

Back to main content list

Scientists are using quantum computing to help them discover signs of life on other planets

Zapata and the University of Hull have partnered to discover how quantum computers could assist with the search for life molecules on exo-planets.

Quantum computers are assisting researchers in scouting the universe in search of life outside of our planet -- and although it's far from certain they'll find actual aliens, the outcomes of the experiment could be almost as exciting. Zapata Computing, which provides quantum software services, has announced a new partnership with the UK's University of Hull, which will see scientists use quantum computing tools to eventually help them detect molecules in outer space that could be precursors to life. Read more at ZD Net here.

Credit: ZD Net

Back to main content list

Coughed particles float for longer in cold air, study suggests

When the temperature drops, turbulent puffs caused by coughs and sneezes become more buoyant and travel further and last longer, scientists in Japan and Italy have discovered.

The researchers say the results of their modelling study could help improve our understanding of the airborne transmission of viruses like SARS-CoV-2. A turbulent puff occurs when a mass of fluid is ejected from a localized source. How they travel and change over time has implications for the dispersal of pollutants, such as those from chimneys, and the transmission of disease droplets in coughs, for example. Read more at Physics World here.

Credit: OIST

Back to main content list

Powering Innovation
Supercomputing in Asia
National Supercomputing Centre (NSCC) Singapore
1 Fusionopolis Way, Connexis South, #17-01 Singapore 138632