./mona.sh_

./university

Monash is a public research university located in Melbourne, Australia. It is Australia's largest University, and provides an education for over 70,000 students.

Monash is a top 100 university, and Australia's leading university for Technology and Engineering. It has a dedicated IT faculty, of which most of the team for the competition are from. This will be the first year that a team from Monash participates in the competition.

./workDistribution

The team has divided up the work in an effective way. Each team member takes the lead on one piece of software, and has a "buddy", who has knowledge about the same application, allowing help when any issues arise.

As we have had meetings throughout the year discussing the applications, we all do have an understanding about all of the applications used in the competition.

./teamMembers

Luhan Cheng, Bachelor of Computer Science
Major: Advanced Computer Science
Focus Application: HPCG

Shuh Fei Chia Bachelor of Computer Science
Major: Advanced Computer Science
Focus Application: Mystery

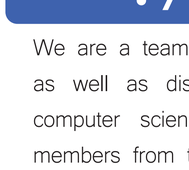

Grace Han, Bachelor of Computer Science
Major: Advanced Computer Science
Focus Application: SeisSol

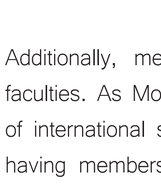

Sam Pardy, Bachelor of Science and Arts
Major: Computational Science and Philosophy
Focus Application: Horovod

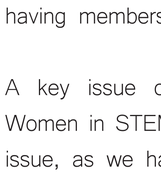

Andrew Persic, Bachelor of Computer Science
Major: Advanced Computer Science
Focus Application: OpenMC

Samuel Stranges, Bachelor of Computer Science
Major: Advanced Computer Science
Focus Application: HPL

./diversity

We are a team that come from a diverse range of backgrounds as well as disciplines. Although most members are studying computer science, we are still educationally diverse, having members from the Arts, Science and IT faculties.

Additionally, members have taken elective courses from other faculties. As Monash is the University with the highest percentage of international students in Australia, we maintain this diversity by having members from multiple countries.

A key issue of diversity within Australia is representation of Women in STEM. We have done our best to alleviate this critical issue, as we have multiple women on the team.

./preparation

Over the course of the year, we have have held weekly meetings in which we discussed various topics relating to the competition; the applications, the hardware, and High Performance Computing fundamentals.

Early on, we had access to CPU cloud boxes provided by The National eResearch Collaboration Tools and Resources project. We were able to practice building the applications from source using different libraries and systems, across a number of different nodes.

In order to maintain the reproducability aspect of the competition, we made it a priority to create build scripts for different architectures, which also allowed quicker implementation on our practice cloud boxes.

Practical, hands on consultations with HPC experts from Australia, from Monash eResearch, the CSIRO and Dell, provided us professional knowlege of the HPC industry, as well as tools and tricks to improve our abilities.

We held several working bees over the course of the year, in which we applied our knowledge installing hardware, configuring a multi node system, scripting, job scheduling and benchmarking.

Application Presentations in front of HPC experts have given us confidence in discussing the applications in preparation for interviews, and deepened our knowldge of the competition applications algorithms and how they can best be optimised.

./powerManagement

A key component of our strategy for power management involved the Dell iDrac. The iDrac is a small, low power consumtion board used to manage the system remotely as well as limit power consumption. One iDrac is installed in each node, allowing us to set a hard limit on power consumption based on hardware power consumption measurements that feed into the iDrac.

Aside from this, nvidia-smi provides monitoring and power limiting capabilities for our GPUs. This tool allows us to run each tesla V100 GPU at as low as 150V. Since the max TDP of our system greatly exceeds the 3KW power limit, and testing has indicated that the V100 GPUs reach peak efficiency at around 180V, we will be relying on nvidia-smi to limit the power consumption of the GPUs so that we stay beneath the power limit while maintaining optimum efficiency.

Lastly, since there is currently no GPU port of openMC, we plan to run openMC on the CPUs in times of low load in order to maximise utilisation of our hardware, which is equipped with 4 Xeon Phi processors.

./software

Operating System:
- Ubuntu Linux 16.04 LTS is used because of its stability, team familiarity, and its powerful package management system.

Accelerator Drivers:
- NVIDIA CUDA 8.0 and 9.0 were chosen for their speed and management tools, as well as compatibility with certain applications for use in the competition.

Compilers:
- GCC and Intel Compilers were chosen. GCC was chosen because of its compatabiliy, while Intel compilers were chosen for the applications that can be built with it for its speed.

Maths Library:
- Intel MKL, chosen for its improved optimisations.

Job Scheduler:
- SLURM was chosen because it is an powerful industry standard tool that is portable and configurable.

Power Management: Dell iDrac is used because of its compatability with the hardware and its powerful and configurable system tools.

./strategies

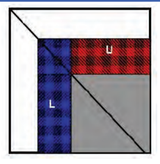

HPL: Using the latest NVIDIA binaries, we have tuned the parameters to better solve the dense system of linear equations on the NVIDIA V100 accelerators.

HPCG: Although there are fewer options when comes to tuning for HPCG, we have optimised our implementation by using an accelerated binary, as well as using the traditional parameter file.

OpenMC: By choosing the best dependencies for OpenMC, we are able to increase speed significanty.

Horovod: We plan on optiminsing Horovod by tweaking the compression of number size, as well as using the NVIDIA Deep Learning SDK.

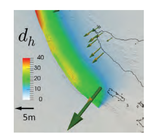

Reproducability: We are able to tune the reproducability for a variety of architectures using the tuning files available to us.

Time Management Strategy: By estimating the time taken for each application, we are able to create a rough timeline for what should be run on which node.

./hardware

The hardware was chosen because of its reliability, efficiency and currency. In total we have two boxes, connected via an Mellanox ConnectX-5 connection. Each box contains the following:

- Motherboard: PowerEdge R940xa
- CPUs: 4 Intel Xeon Gold 6154
- Volitile Memory: 192GB of Dell 1RX8 DDR4 RAM
- Accelerators: 4 NVIDIA V100 Cards
- SATA Storage: Dell 480GB Midex Use SSD
- NVME Storage: Dell 256 M.2 PCIe 2230-S2

./whyWeWillWin

We believe that we have the ability to win, because we have the following:

- Faculty support
- Knowledge and deep understanding
- Powerful hardware
- Skills
- Significant preparation
- Winning strategies

./sponsors

NVIDIA

DELL

CSIRO

ANCORA IMPARO