GPUs on NeSI

Chris Scott Wolfgang Hayek Alex Pletzer Maxime Rio Georgina Rae

New Zealand eScience Infrastructure (NeSI)

contact: chris.scott@nesi.org.nz

11 February 2021

Outline

Background

CPUs and GPUs

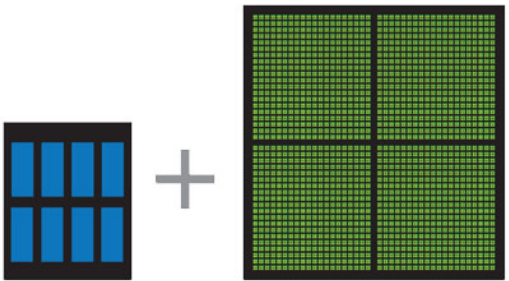

- GPUs (Graphics Processing Units) are accelerators that can be used with CPUs to boost the performance of many applications.
- Offload compute intensive portions of the code to GPUs and leave the remainder of the code on the CPU
- Design philosophy
  - CPUs consist of a small number of powerful cores
  - GPUs consist of thousands of lighter weight cores, designed to process parallel workloads efficiently

CPU vs GPU

CPU vs GPU

CPU vs GPU

GPUs: why now?

- Current GPUs on NeSI
  - ~8 NVIDIA P100s on Mahuika (some more on Maui Ancillary nodes too)
- Adding new GPU capability
  - NVIDIA A100 and AMD MI100
  - Should be a significant improvement (more powerful, more memory, . . .)
  - Coming soon . . .
- Good time to think about using GPUs
  - Performance is increasing significantly
  - Getting easier to use

How can you utilise GPUs?

1. Your software already has support for GPUs or has been ported by somebody else
   - May be as simple as requesting a GPU from Slurm (low effort)
   - Good support in high level languages: Python cupy, MATLAB gpu arrays, Julia, ...
2. Calling GPU libraries (cuBLAS, cuFFT, ...) to offload expensive calculations
   - Often low effort (especially if the GPU library has the same API as non-GPU; could just require relinking)
3. Adapt your code to offload loops onto the GPU using an API like OpenACC
   - Some more effort required but generally doesn't require big changes to a code base; could just be adding pragma statements to loops
   - Good support from compilers is limited (PGI, maybe Cray and GNU)
4. Writing custom GPU kernels using a language such as CUDA
   - More effort and could require significant changes to code base (more maintenance, possibly less portable) but greatest flexibility

GPU case studies

Case study – GPU support in the software

Protein Modelling – Dr Wanting Jiao, VUW

- Why GPUs?
  - Very computationally expensive
  - Know that NAMD come with good GPU support builtin
- How did we utilise GPUs?
  - NVIDIA provide a container for NAMD: `https://ngc.nvidia.com`
  - Very simple to run on the GPUs: `https://support.nesi.org.nz/hc/en-gb/articles/360001500156-NVIDIA-GPU-Containers`
- Outcome
  - Performance on 1 P100 GPU is roughly the same as 3 full Maui nodes (120 cores)

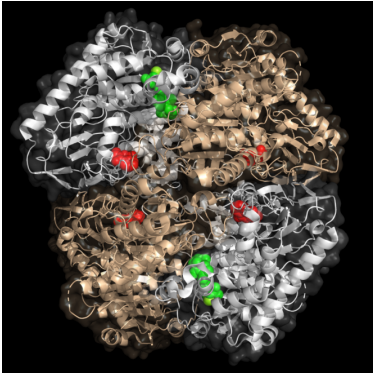

Case study – GPU libraries

Tropical Circulation Model – Dr Gilles Bellon, UoA

- Why GPUs?
  - Wants to move to higher resolution grids
  - Matrix multiplication is a bottleneck
- How did we utilise GPUs?
  - MKL for multithreaded dgemm
  - Added an option to use dgemm from cuBLAS
  - CMake enables easy switching
- Outcome
  - GPU matrix multiplication 36% faster than MKL 16 threads and comparable to MKL 32 threads

Case study – OpenACC

High Performance Marketing Insights – Dr Damien Mather, UoO

- The bottleneck in his current approach is the log determinant calculation, $\mathcal{O}(n^3)$
- Started with an MPI implementation of the Condensation method
- Adding a couple of lines of OpenACC directives gave a 13x speedup over the serial CPU version
- Managed memory makes it really easy to try offloading loops to the GPU (no need to explicitly copy data between host and device)

Case study – OpenACC

High Performance Marketing Insights – Dr Damien Mather, UoO

- NeSI Consultancy to optimise data locality
  - Additional 2.1x speedup
  - Increased code complexity
- P100s perform similarly to full Maui nodes
- Most of the gains came from adding a couple of lines of OpenACC directives!
- Not invasive – compiling without OpenACC flags still works

Case study – Custom CUDA code

Solar System – Dr Philip Sharp, UoA

- Based on the CUDA N-body reference implementation, $O(N^2)$
- Thousands of bodies maps to thousands of threads on the GPU
- Each thread computes (part of) the acceleration for a single body
- This calculation is not feasible on a CPU only

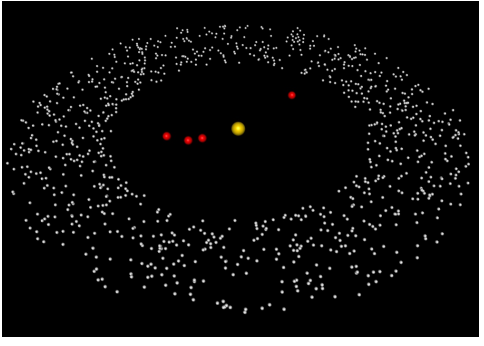

Case study – Custom CUDA code

Solar System – Dr Philip Sharp, UoA

- NeSI Consultancy to implement a multi-GPU version (within a single node) to squeeze out even more performance
- Use CUDA pinned memory to get truly asynchronous memory copies between a single host thread to multiple devices
- 1.7x speedup for numbers of bodies – good boost for an already optimised code

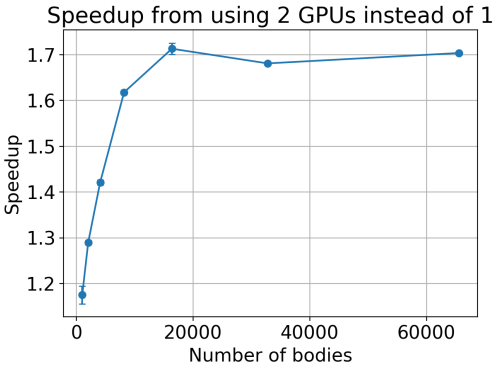

Summary

Summary

- NeSI is getting new GPUs
- Good time to think about whether you can take advantage of them
  - Performance and ease of use have increased a lot
  - Does not always require much time or effort
- If your code doesn't support GPUs already then OpenACC or GPU libraries (such as cuBLAS) are a good place to start looking
- Always profile first to make sure you aren't wasting your time
- NeSI Consultancy can help:
  `https://www.nesi.org.nz/services/consultancy`

NeSI New Zealand eScience Infrastructure

RSE RSE-AUNZ

Science New Zealand

NZ Research Software Engineers Conference

Spring 2021

Help us plan the programme!
Email events@nesi.org.nz to get involved.

Who attends:

- Researchers and academics who code
- Software engineers & system admins working in the research domain
- Generalists who bring together the research and technical domains
- Developers, IT managers, coding enthusiasts, and big data analysts from Crown Research Institutes, universities, and other public sector organisations

NeSI @ eResearch NZ - Talks & Workshops:

Wednesday 10 Feb

13:00 - 17:00 - **Maxime Rio -** Machine Learning on NeSI 101

13:20 - 13:40 - **Jun Huh -** Taonga: building a data repository for genomics research in New Zealand

13:20 - 13:40 - **Dinindu Senanayake -** Paving the way for Bioinformatics excellence in New Zealand

14:20 -15:00 - **Brian Flaherty** - Moving data: getting up to speed with Globus and Science DMZ

15:50 - 16:50 - **Jana Makar** - Challenge Accepted: Responding to community feedback for supporting diversity in HPC & eResearch

Thursday 11 Feb

11:00 - 11:20 - **Maxime Rio** - Data science consultancies at NeSI: A whirlwind tour of case studies

13:30 - 13:50 - **Chris Scott** - GPUs on NeSI

13:50 - 14:10 - **Georgina Rae** - Building Partnerships for eResearch

14:10 - 14:30 - **Wolfgang Hayek** - NeSI Consultancies - Evolving a Scientific Programming Service

14:40 - 15:00 - **Albert Savary** - Software on NeSI

15:00 - 15:20 - **Jeff Zais** - Taking Advantage of Technology Innovations in the Next Generation of NeSI HPC Infrastructure

15:20 - 15:40 - **Callum Walley** - Virtual Desktops for HPC

Thursday 11 Feb (cont.)

15:20 - 15:40 - **Robin Bensley** - Staying connected in an evolving eResearch ecosystem

16:00 - 17:00 - **Megan Guidry** - Sowing the seeds of capability: Experience what Carpentries instructor training is all about

Friday 12 Feb

11:20 - 12:30 - **Nick Jones** - Future of eResearch

12:20 - 12:30 - **José Filipe Gonçalves Higino** - Coaching great practices of describing a problem

13:30 - 14:30 - **Blair Bethwaite** - Embracing cloud-native architectures

13:30 - 14:30 - **Alexander Pletzer and Nooriyah Lohani** - Who needs research software engineers?

13:30 - 14:30 - **Georgina Rae** - FAIR for Research Software