

HPC Workshop @ RSE18

Birmingham 5th September 2018

Topic: *Benchmarking*

Group leader: Julien Highfield

Scribe: Andy Turner

Reporter: Mark Williamson

Other participants

- Tomas
- Chris Hadjigeorgiou
- Neelofer Banlawala
- Feng Chen

Notes from discussion

- Why did you come?
 - About to run a procurement
 - Optimising for our Tier2 centre - what is good performance? How can we help people make best use of the system
 - Understand how to do benchmarking. Are there any tools available for benchmarking and profiling? Which should I pick and why? Where do I get simple examples on how to do this?
 - This also came back from ARCHER User Survey: can you provide a howto on how to do this?
 - Julian did start to write simple howto's on this but did not match the community
 - Opportunity to provide this usefully for the community both users and RSEs
 - Good practical way to get an idea of differences between architectures
 - Statistical benchmarking - what is the performance difference. How many benchmark runs do I need to understand the statistics. Especially with JIT compiling etc.
 - What is the distribution of performance for this benchmark. Measure by running periodically at randomised times
 - Can we use test
 - Want to understand if we can have a common set of benchmarks and approach across different sites
 - Can I contribute to the open repositories? If so, how?

- What is a fair comparison?
 - For example, LAMMPS can include the user-intel update that improves the performance of the code on Intel. Is this fair on other architectures?
- What should be in a benchmarking repository?
 - How broad should it be?
 - Should ML be included? Should Python be included?
 - What size of benchmark should be included? Benchmarks age quite quickly?
 - What is the strategy for updating over time?
- JMH is a Java microbenchmark harness that allows you to annotate the code
- JUBE from Juelich can automate benchmarking
 - Experience from a couple of people that it was quite difficult to use
- When procuring, how much are vendors allowed to modify codes?
- How do you make it worth peoples' while to contribute back to the repository?
 - Need clear guidelines on what can be contributed and what you need
- How do you deal with licensed material when sharing benchmarks and results?

Summary

- General appetite for more information on how to benchmark and profile and what tools are available to help?
 - Very useful for RSEs and beyond working on HPC
- Benchmark reproducibility and how to measure, how to set it up? Techniques and strategies
- What should be in a common repository?
 - How do you moderate content?
 - How do you maintain/update benchmarks?
- Scope for pulling together a UK benchmarking interest group