European Centre of Excellence in Physics at Extreme Scales

Higgs Centre for Theoretical Physics
April 15, 2014

Tuesday, 15 April **Higgs Centre (Room 4305)**

08:30–09:00 Registration and tea/coffee/light snack

09:00–09:10 Objectives and plan for the day
Speaker: Richard Kenway (Edinburgh)

09:10–09:30 Horizon 2020 and PRACE
Speaker: Alison Kennedy (Edinburgh)

09:30–09:45 Discussion

09:45–10:00 Precision hadron structure & novel computer architectures
Speaker: Constantia Alexandrou (The Cyprus Institute)

10:00–10:15 Cambridge contribution to the Centre of Excellence
Speaker: Matthew Wingate (Cambridge)

10:15–10:30 Phi in the Sky and the Extreme Universe
Speaker: Paul Shellard (Cambridge)

Abstract: I will describe the innovative developmental work currently being undertaken on the COSMOS@DiRAC supercomputer, a hybrid 15TB shared-memory/MIC co-processor system. This work relates to physics at the most extreme energies of all, modelling predictions from the inflationary phase of the Big Bang and testing these against observational data from the cosmic microwave sky and billion galaxy surveys. I will discuss the benefits and pathways forward for flexible hybrid systems and the programming paradigms they support.

10:30–10:45 Supercomputer simulations of galaxy formation
Speaker: Tom Theuns (Durham)

Abstract: I will discuss the Eagle simulations that follow the formation and evolution of a cosmologically representative set of galaxies. I will further discuss recent work on developing a task-based parallel code called Swift designed to improve load balance when performing Eagle-like simulations on mixed architectures.

10:45–11:00 Co-design for lattice QCD
Speaker: Peter Boyle (Edinburgh)

11:00–11:15 Tea and coffee

11:15–11:30 Automated Code Generation for Lattice QCD Simulation
Speaker: Denis Barthou (Inria)

Abstract: Quantum Chromodynamics (QCD) is the theory of strong nuclear force, responsible of the interactions between sub-nuclear particles. QCD simulations are typically performed through the lattice gauge theory approach, which provides a discrete analytical formalism called LQCD (Lattice Quantum Chromodynamics). LQCD simulations usually involve generating and then processing data on petabyte scale which demands multiple teraflop-years on supercomputers. Large parts of both, generation and analysis, can be reduced to the inversion of an extremely large matrix, the so-called Wilson-Dirac operator. For this purpose, and because this matrix is always sparse and struc- tured, iterative methods are definitely considered. Therefore, the procedure of the application of this operator, resulting in a vector- matrix product, appears as a critical computation kernel that should be optimized as much as possible. Evaluating the Wilson- Dirac operator involves symmetric stencil calculations where each node has 8 neighbors. Such configuration is really hindering when it comes to memory accesses and data exchanges among processors. For current and future generation of supercomputers the hierarchical memory structure makes it next to impossible for a physicist to write an efficient code. Addressing these issues in other to harvest an acceptable amount of computing cycles for the real need, which means reaching a good level of efficiency, is the main concern of this talk. We present here a Domain Specific Language and corresponding toolkit, called QIRAL, which is a complete solution from symbolic notation to simulation code.

11:30–11:45 Programming Heterogeneous Computing Architectures with the StarPU Runtime System
Speaker: Olivier Aumage (Inria)

Abstract: Heterogeneous computing architectures such as multicores equipped with accelerators are notoriously difficult to program due to the strong differences in performance characteristics among the various available computing units and also to the discrete memory spaces of accelerating boards.

In line with proposed CoE WP3 scope, the StarPU runtime system developed by INRIA Runtime Team in Bordeaux aims at providing abstraction and portability of performance to application programmers. It dynamically maps elementary computing tasks on available computing units in such a way as to maximise the overall execution efficiency, and it transparently data movements. Basically, the application provides algorithms (CPU/GPU implementations of tasks) and constraints (task dependencies). Then, StarPU handles run-time concerns: Task dependencies enforcement; Optimised heterogeneous scheduling; Optimised data transfers and replication. Moreover, StarPU provides an OpenCL compatibility interface that allows OpenCL-enabled applications to be readily ported on top of StarPU with little effort.

11:45–12:00 Programming models and compiler support for heterogeneous parallelism
Speaker: Andrew Richards (Codeplay)

Abstract: I will provide a brief overview of Codeplay's background in C++ and OpenCL compiler R&D, and performance optimisation. On behalf of the Khronos subgroup, SYCL, I will then describe both the provisional SYCL specification, which allows a portable, single-source C++ layer to access OpenCL hardware; and our implementation of it.

12:00–12:15 GPUs for Physics
Speaker: Alan Gray (Edinburgh)

12:15–12:30 Developing code for Petascale and Exascale hybrid environments
Speaker: Avtar Cheema (Allinea Software)

Abstract: The HPC market is rapidly changing with the adoption of a growing number of cores per system and the adoption of new emerging technologies such as accelerators. This change is increasing the pressure on the Researchers and Scientists who need to port/revamp their application(s) so that they can fully exploit the new hardware capabilities. The question placed before us now is 'Are the tools currently on the market able to support the users in completing their task at petascale and exascale levels?'

For several years, we at Allinea have been collaborating with multiple large organizations in EMEA and in the USA to help address these needs. Within the collaborations, when we originally began in 2003 our quest was to develop Allinea DDT, so as to enable at-scale debugging. In this we found that two of the key industry requirements for a parallel debugger were efficiency and responsiveness. However, we quickly realized that this was only a milestone. What also needed to be urgently addressed was the challenge of how to improve the way in which data is presented to developers regardless of the size or scale of the debugging session; a suitable GUI was essential.

Today, we are pleased to say that Allinea develop and provide our proprietary software development tools:

- Allinea DDT - the most widely used parallel debugger
- Allinea MAP - an excellent companion capable of profiling without intrusion
- Allinea Performance Report - HTML report for parallel app performance analysis

These tools have been designed with the following points in mind:

- User community – a strong self-servicing user community who can work independently of the HPC System Administrator
- Efficiency – A unique intelligence that drastically reduces the developers time to debug and profile complex parallel codes.
- At Scale debugging - the ability to work at Scale and successfully present this data to the users in a succinct manner
- Easy to use - a HTML application performance analysis report and GUIs that are intuitive and very easy to navigate through, resulting in less time on the cluster and more time researching.

12:30–12:45 Computing for the LHC
Speaker: Pete Clarke (Edinburgh)

12:45–13:15 Lunch

Tuesday, 15 April **Higgs Centre (Room 4305)**

13:15–13:30 Snippets from the DiRAC-3 Science Case
Speaker: Jeremy Yates (STFC DiRAC)

13:30–13:45 Training Europe wide
Speaker: David Henty (Edinburgh)

13:45–14:30 Discussion on scope and structure of the CoE

14:30–15:00 Assignment of Work Package leaders and participants

15:00 Tea and coffee

15:00–15:30 Parallel session for Work Package groups to discuss plans

15:30–15:45 Work Package leaders report

15:45–16:00 Next steps
Speaker: Richard Kenway (Edinburgh)

16:00 Close