

Prof. Jan HESTHAVEN

Mathematics Institute of Computational Science and Engineering - MATHICSE

SEMINAR OF NUMERICAL ANALYSIS

➤ **TUESDAY 6 OCTOBER 2015 - ROOM GC D0 386 - 16h15**

Prof. J. TROMP, (Princeton University, USA) will present a seminar entitled:

"Towards exascale seismic imaging & inversion: taming workflow and I/O issues"

Abstract:

Providing a better understanding of the physics and chemistry of Earth's interior through numerical simulations has always required tremendous computational resources. Post-petascale supercomputers are now available to solve complex scientific problems that were thought unreachable a few decades ago. They also bring a cohort of concerns tied to obtaining optimum performance. Several issues are currently being investigated by the HPC community. These include energy consumption, fault resilience, scalability of the current parallel paradigms, workflow management, I/O performance and feature extraction with large datasets. In this presentation, we focus on the last three issues.

In the context of seismic imaging -in particular for simulations based on adjoint methods- workflows are well defined. They consist of a few collective steps (e.g., mesh generation or model updates) and of a large number of independent steps (e.g., forward and adjoint simulations of each seismic event, pre- and postprocessing of seismic traces). The greater goal is to reduce the time to solution, that is, obtaining a more precise representation of the subsurface as fast as possible. This brings us to consider both the workflow in its entirety and the parts comprising it. The usual approach is to speedup the purely computational parts based on code optimization in order to reach higher FLOPS and better memory management. This still remains an important concern, but larger scale experiments show that the imaging workflow suffers from severe I/O bottlenecks. Such limitations occur both for purely computational data and seismic time series. The latter are dealt with by the introduction of a new Adaptable Seismic Data Format (ASDF). Parallel I/O libraries, namely HDF5 and ADIOS, are used to drastically reduce the cost of disk access. Parallel visualization tools, such as VisIt, are able to take advantage of the ADIOS metadata to extract features and display massive datasets. Because large parts of the workflow are embarrassingly parallel, we are also investigating the possibility of automating the imaging process with the integration of scientific workflow management tools, such as Pegasus, Kepler, or Swift.

Lausanne, 9 September 2015/JH/cr