Geophysical Research Abstracts
Vol. 15, EGU2013-4483, 2013
EGU General Assembly 2013

VERCE - CPU-intensive Applications in Seismology

Marek Simon (1), Siew Hoon Leong (2), Kasra Hosseini Zad (1), Lion Krischer (1), Michele Carpene (3), Graziella Ferini (3), Luca Trani (4), Alessandro Spinuso (4), Federika Magnoni (5), Emanuele Casarotti (5), André Gemünd (6), David Weissenbach (7), Iraklis Klampanos (8), Heiner Igel (1), and the VERCE (www.verce.eu) Team

(1) Earth and Environmental Sciences, Ludwig-Maximilians-University Munich, Munich, Germany, (2) Leibniz-Rechenzentrum (LRZ) der Bayerischen Akademie der Wissenschaften, Garching, Germany, (3) Cineca, Bologna, Italy, (4) Koninklijk Nederlands Meteorologisch Instituut (KNMI), DeBilt, Netherlands, (5) Instituto Nazionale di Geofisica e Volcanologia (INGV), Roma, Italy, (6) Fraunhofer-Institut für Algorithmen und Wissenschaftliches Rechnen (SCAI), Sankt Augustin , Germany, (7) Institut de Physique du Globe de Paris (IPGP), Paris, France, (8) School of Informatics, University of Edinborough, Edinborough, UK

Recently, advances in computational seismology have culminated in the development of a range of scientific codes enabling the calculation of highly accurate 3D wave and rupture propagation in complex 3D media at unprecedented scales and level of detail. Fortunately, the computational hardware has grown at rates at least as vigorous, to match up to the heavy requirements in CPU and memory imposed by realistic applications. However, as algorithmic and hardware complexity increases, making them work efficiently has become difficult: legacy codes need to be adapted and maintained by the community to meet the requirements of the new computational environments and the handling of large volumes of expensively generated data has become a challenge in itself.

Within the VERCE (www.verce.eu) project, several specific use cases have been developed, exemplifying the challenges ahead. Seismic 3D-forward modelling of a large number of recorded earthquakes on a continental scale represents a model use case involving HPC. The simulation will be carried out on an HPC machine (SuperMUC, PLX), the resulting data submitted to a publicly accessible community Data-Center (ORFEUS) with the possibility to interactively mine and process the data using Grid infrastructure (Fraunhofer-SCAI, IPGP). As this basic workflow will need to be repeated for each solver, model, frequency range or processing option over and over again, the elements need to be connected within a workflow environment, allowing easy customization, job monitoring and visualisation of results.

In collaboration with our VERCE partners, it was possible to define a basic core architecture for the VERCE platform for the proposed use case. Currently established components include JSAGA for job submission to GRAM, gLite Cream, gLite WMS as well as UNICORE6 instances, GridFTP for file transfer, using VOMS enabled certificate-based authentification. Additionally, a few suggested community applications (Seissol, Specfem3D Sesame, Ses3d, Axisem) as well as tools and libraries (e.g. Obspy for data processing) have been installed on the relevant sites complying to the standards set by the involved organisation (EGI, PRACE).

Due to the paradigm of "separation of concern" there will be different levels at which a user might interact with the platform. Besides the possibility for "advanced users" of using the platform as a model framework for different use cases, it is anticipated that the platform will feature an interactive browser-interface that will allow easy initialisation of HPC jobs while allowing seamless access to visual representation of the results. In the presentations we will highlight recent advances and point out remaining challenges to generate community interest in the scientific platform.