

Press release
For release: 16/11/2009

PRACE is Ready for Implementation: Applications ported

PRACE, the Partnership for Advanced Computing in Europe, has explored promising petascaling techniques. Related work has been done on optimization techniques and on studying software libraries and programming models suitable for petascaling. This work has prepared the way for the efficient exploitation of the upcoming Tier-0 systems.

The applications studied in this work cover a broad range of scientific areas, and are representative of the European HPC usage. Most of them also originate from the European scientific community. The applications are: Alya, AVBP, BSIT, Code_Saturn, CP2K, CPMD, Echam5, Elmer, EUTERPE, Gadget, GPAW, Gromacs, HELIUM, NAMD, NEMO, NS3D, Octopus, PEPC, SIESTA, SPECFEM3D, QCD, Quantum_Espresso and WRF. These applications were ported, evaluated and scaled on the PRACE prototypes, which represent the current top of the line supercomputer architectures.

The applications were ported, on average, to three prototype systems. As was to be expected, the difficulty of porting an application to a prototype depended to a large degree on how "exotic" the prototype is. The cluster based systems posed perhaps the least problems in porting, while the few codes that were ported to the Cell based prototype required a significant investment in time. Some problems were encountered and solved, such as limited library support, cross compilation issues and compilers producing non-working code at higher optimisation levels.

As for optimization, it proved essential to tune the options and parameters used when compiling and running a code, such as the choice of numerical libraries and compiler options. A tool was developed for studying optimal compiler options and platform-specific recommendations were produced. For many codes there was limited gain in optimizing the performance of the computational parts, as the performance was limited by the efficiency of external numerical libraries and the scalability of the application. The scalability was improved using a variety of techniques, for example, parallel I/O, improved collectives and hybrid programming techniques. Best practices were collected by PRACE and the application evaluation documents report these. They also show the suitability of a given architecture for running a particular code.

PRACE also investigated the programming models and the software libraries required by petascaling applications. An accurate survey and analysis of the new upcoming programming models and languages suitable for petascale applications has been completed. In summary, MPI is still the *de facto* standard for HPC, but fundamental limitations are critical for hyper

www.prace-project.eu

scalability. A more promising approach, to effectively scale up the applications, seems the hybrid programming model, based on coupling message-passing with multi-threading. Surprisingly few of the applications employ this technique fully at the moment. However, many are planning to do so in the future.

Different players are proposing new parallel languages, but none of them appears to be able to reach a production level maturity in the short term, and have a strong impact on the computational community. The PGAS languages could be promising mainly if a standardisation process will be effective soon.

Hardware accelerators appear to have a big potential. At present time it is difficult to identify common directions and guidelines for languages to efficiently exploit accelerators. The language which delivers best performance is CUDA, but most people do not want to be limited to one kind of accelerator and pin their hopes on OpenCL instead. Another important issue in the direction of portability and productivity are tools like RapidMind and CAPS HMPP, but it is still too early to see what their real impact will be. Probably a good direction to enhance the programmability for accelerators could be the availability of compilers integrating directly specific APIs to allow a semi-automatic code generation.

The work carried out by PRACE has benefited a number of European HPC applications. The applications have been ported and optimised to the new PRACE prototype petascale systems, and best practices on this work have been collected. This lays the groundwork for reaching petaflop/s performance on the future PRACE Tier-0 petascale machines.

More information: Sebastian von Alfthan, e-mail: sebastian.von.alfthan@csc.fi, Giovanni Erbacci, e-mail: g.erbacci@cineca.it, Mohammad Jowkar, e-mail: mohammad.jowkar@bsc.es

About PRACE: The Partnership for Advanced Computing in Europe prepares the creation of a persistent pan-European HPC service, consisting of several tier-0 centres providing European researchers with access to capability computers and forming the top level of the European HPC ecosystem. PRACE is a project funded in part by the EU's 7th Framework Programme (FP7/2007-2013) under grant agreement n° RI-211528.