TERRA software

Introduction

The TERRA code of Dr Huw Davies is used on an HPC Wales project HPCW045. It performs simulations of the Earth's plate tectonics which result from mantle flow beneath the earth's crust to underpin a better understanding of earthquakes and volcanic eruptions, which generally take place on plate boundaries.

Dr Davies had requested support from HPC Wales in order to assess the performance of the software and what enhancements could be added to improve its execution on the Sandy Bridge nodes. In addition they requested help in integrating work from a different development path back into their main source code repository.

Overall, it was clear that the optimisation work performed on the TERRA code showed that it has been well implemented and developed over the years. In order to ensure the long-term portability of the software and aiding the compiler in getting the best performance out of the code, as well as helping future developers, a list of pointers for current software engineering best practice were provided to Dr Davies and his team, many of which we believe have since been implemented.

Parallel profiling

An initial profiling of the code was conducted using the packages Tau and Scalasca, both of which are installed on the HPC Wales resources. This showed that almost all the computational work was spent in a very small number of routines. The MPI communications were shown not to be a significant overhead. It was found that there would be possibilities for small performance gains, but generally the code was performing very well.

When performing the optimisation work it was seen how a few percent of extra performance could be extracted out of the main computationally intensive routines. However, as expected, the performance per process was as generally good as could be expected. This was evidenced though a combination of parallel profiling and examination of the memory access patterns in these routines. The difficulties in enhancing performance much further were due to the memory access the large tensor arrays necessary from the mathematics of the problem being solved, implemented as six levels of nested loops. This meant that the performance of increasing vectorization, etc., was of no benefit since the limiting factor was that there could be no significant reuse of data in cache.

Ewrop & Chymru: Buddsoddi yn eich dyfodol
Cronfa Datblygu Rhanbarthol Ewrop
Europe & Wales: Investing in your future
European Regional Development Fund
ERDF
Llywodraeth Cymru
Welsh Government

Branch re-integration

In 2011, a version of the code was branched off from the community version by John Baumgardner. Although some of this version is now outdated, it does include what is believed to be a much better method of solving the energy equation. The strategy employed to do this in the current community version of TERRA can produce instabilities especially where there are sharp temperature gradients in the temperature field, yielding an oscillatory solution. This can sometimes feed back into the whole system. The alternative approach implements a semi-Lagrangian method for the heat transport that appears to be extremely stable. The current version of TERRA uses a Runge-Kutta scheme to update the temperature and so the intention for re-integration, was to replace this with the semi-Lagrangian method from the 2011 Baumgardner branch.

As a first step to understanding the changes made by Baumgardner to the code, it was necessary to establish the version of the code from when it had been branched. Digital archaeology revealed that the branch was taken in April 2011, source code version 699.

Overall there were very significant changes to the code between all three versions (trunk, r699 and the Baumgardner version), but it has been possible to put in the semi-Lagrangian formulation into the latest version of the main code. The main issues with this integration were concerned with matching the COMMON blocks and arrays between the codes.

The time integration method from the Baumgardner version was integrated successfully enough for Dr Davies and his team to complete the final debugging of the results in order to verify how well the new method matched the original method they used as well as the results for simplified cases that may be run using the pure Baumgardner branch. Unfortunately Dr. Davies' team were unable to get this new version to match benchmark results in detail across a wide number of cases. Further work is needed before this advance can be taken forward into the production code.