Full Stokes finite-element modeling of ice sheets using a graphics processing unit

Hakime Seddik[1], Ralf Greve[1]

[1]*Institute of Low Temperature Science, Hokkaido University, Japan*

Thermo-mechanical simulation of ice sheets is an important approach to understand and predict their evolution in a changing climate. For that purpose, higher order (e.g., ISSM, BISICLES) and full Stokes (e.g., Elmer/Ice, http://elmerice.elmerfem.org) models are increasingly used to more accurately model the flow of entire ice sheets. In parallel to this development, the rapidly improving performance and capabilities of Graphics Processing Units (GPUs) allows to efficiently offload more calculations of complex and computationally demanding problems on those devices (Cecka and others, 2011; Markall and others, 2013). Thus, in order to continue the trend of using full Stokes models with greater resolutions, using GPUs should be considered for the implementation of ice sheet models. We developed the GPU-accelerated ice-sheet model Sainō. Sainō is an Elmer (http://www.csc.fi/english/pages/elmer) derivative implemented in Objective-C which solves the full Stokes equations with the finite element method. It uses the standard OpenCL language (http://www.khronos.org/opencl/) to offload the assembly of the finite element matrix on the GPU. A mesh-coloring scheme (Komatitsch and others, 2010) is used so that elements with the same color (non-sharing nodes) are assembled in parallel on the GPU without the need for synchronization primitives. The current implementation shows that, for the ISMIP-HOM experiment A, during the matrix assembly in double precision with 8000, 87,500, 252,000 and 391,500 brick elements, Sainō is respectively 7x, 17x, 23x and 25x faster than Elmer/Ice (when both models are run on a single processing unit). In single precision, Sainō is even 9x, 38x, 45x and 49x faster than Elmer/Ice. A detailed description of the comparative results between Sainō and Elmer/Ice will be presented, and further perspectives in optimization and the limitations of the current implementation.

References

Cecka, C., Lew, A.J. and E. Darve, Assembly of finite element methods on graphics processors, Int. J. Numer. Meth. Engng, 85, 640-669, 2011.

Komatitsch, D., Erlebacher, G., Göddeke, D. and D. Michéa, High-order finite-element seismic wave propagation modeling with MPI on a large GPU cluster, Journal of Computational Physics, 229(20), 7692-7714 , 2010.

Markall, G. R., Slemmer, A., Ham, D. A., Kelly, P.H.J., Cantwell, C.D. and S.J. Sherwin, Finite element assembly strategies on multi-core and many-core architectures, Int. J. Numer. Meth. Fluids, 71, 80-97, 2013.