Accelerating 2D and 3D Ocean models with the Discontinuous Galerkin Method and GPU Computing

MARINE 2023

Miguel De Le Court*, Vincent Legat and Jonathan Lambrechts

Institute of Mechanics, Materials and Civil Engineering (IMMC)
UCLouvain
Louvain-la-Neuve, Belgium
e-mail: miguel.delecourt@uclouvain.be, web page : https://www.slim-ocean.be/

ABSTRACT

With a growing number of supercomputers getting most of their processing power through the use of graphics processing unit (GPUs), developing software capable of using them is becoming a necessity. In this work, we develop a massively parallel version of the SLIM [1] ocean model optimized to run on multiple GPUs or CPUs.

Our model uses the Discontinuous Galerkin Finite Elements method (DG-FEM) on non-structured 2D meshes extruded vertically to form superimposed layers of prisms. The algorithms for those are adapted and optimized to map to the Single Instruction Multiple Threads (SIMT) architecture of GPUs. Due to the unstructured nature of the mesh, special care is put into the memory access pattern, with significant changes required in order to maximize coalescence and locality. A domain decomposition strategy is also implemented to scale out to a large number of GPUs.

The performance and accuracy of our model is then validated with relevant benchmarks and case studies. We show that our model achieves between 20% and 50% of the peak floating-point performance of GPUs, thus making a single GPU comparable to more than 500 CPU cores although only using a fraction of the power. While our implementation is mainly optimized for GPUs from NVIDIA and AMD, we also take a step back and look at the broader state of the silicon industry to see what design lessons scientific software can learn from GPU programming.

REFERENCES

[1]   E. Deleersnijder *et al.*, "SLIM: A multi-scale model of the land-sea continuum," presented at the VLIZ Marine Science Day 2019, 2019. Accessed: Feb. 13, 2022. [Online]. Available: https://dial.uclouvain.be/pr/boreal/object/boreal:243281