Parallelizing a shallow-water flow model with variable pressure

Executive Summary

A team from the University of Malaga, Spain, attended the second Appentra-led hackathon, CESGAHACK-18 and made use of Appentra's unique parallelization methodology to parallelize their shallow-water model that models salinity concentrations and temperature to understand natural flows. The model uses a computational approach that is faster than the direct simulation of flows using Navier-Stokes equations while providing accurate simulations. It has the potential to be used in a wide range of geophysical flow applications.

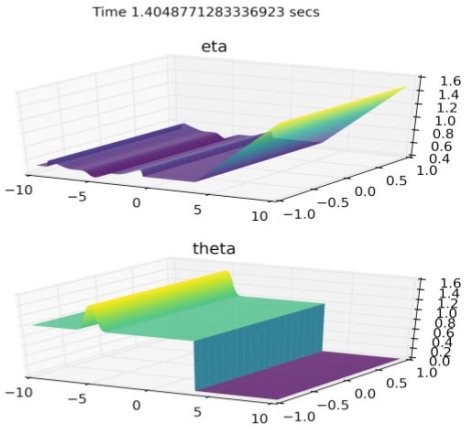

The Challenge

The simulation of oceanographical currents and stratified atmospheric flows is computationally time-consuming and technically challenging. Shallow water simulation models could be used to understand how deep waters evolve and how they could be affected by external factors such as climate change. However, problems arise in stratified atmospheric flows where a more sophisticated approach must be taken in order to capture reliably the vertical behavior of the flow. The model presented is able to accurately simulate complex geophysical flows in a cost-efficient manner compared with traditional fluid-flow simulation methods, enabling a clearer understanding in a shorter amount of time of ocean currents and atmospheric flows. To further reduce the computational time the use of HPC is essential.

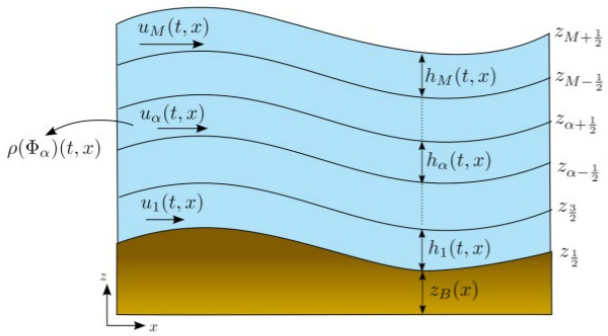

Figure 1: Sketch of the multilayer approach

"It was a great organization for a great event where everything worked, from the advisors to the participants. We would have had a hard time achieving this level of parallelism in just one week in our own. We are very happy with the experience and we are looking forward to the next one!"

Ernesto Guerrero
University of Malaga
EDANYA Group

Parallelizing the Model

During CESGAHACK-18 the EDANYA team successfully parallelized their code. They produced two working versions: OpenMP and OpenACC enabled codes, providing the team with software that works on both shared memory architectures and GPGPUs and that is highly portable across multiple HPC services. Within the 5 days of the hackathon they first added OpenMP capabilities, which provided a 15 times speedup on 32 cores for 160,000 volumes.

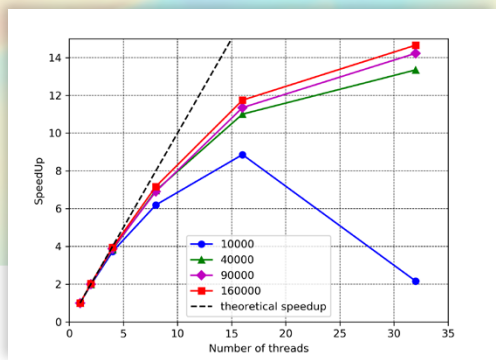

They also added OpenACC capabilities and achieved a 47.79 times speedup for 160,000 volumes compared to their serial code. With OpenACC the team increased the precision of the numerical solution up to 1,440,000 volumes and achieved a 63.97 speedup. Crucially, their experience has provided them with a set of methodologies to continue improving the code and overall performance.

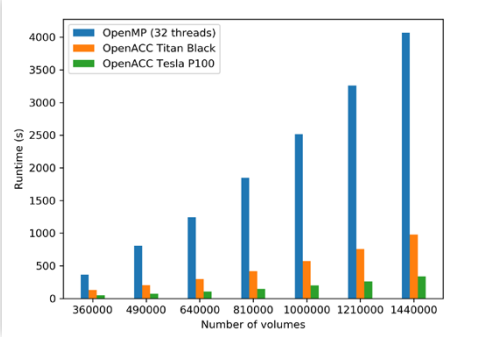

We gratefully thank to the Supercomputing Centre of Galicia (CESGA) for supporting the organization of the GPU Hackathon and for providing access to the FinisTerrae supercomputer.