EAGE

HPC10

Better Productivity and Portable Finite Difference Wave Equation Propagators Using Directive Based Programming

G. Hugues* (Total E&P Research & Technology), H. Calandra (Total)

Summary

Heterogeneous architectures such as GPU have demonstrated during the past decade to be highly efficient and provide better performance than conventional CPU for seismic depth imaging. GPUs requires the use of specific language extension that are considered as low-level programming model such as CUDA. Unfortunately, CUDA is not standard and not portable on multiple vendors and hardware. OpenACC is standard high-level programming model based on directives that aims to simplify GPU programming. Despite offering a more productive way to develop parallel code on GPU, OpenACC has been considered way less efficient than low-level programming models like CUDA or OpenCL. In order to understand what is the performance gap between those models, we have implemented finite difference wave equation propagators using OpenACC. In this paper, we are demonstrating that OpenACC can achieve up to 70% of efficiency when compared to the best CUDA implementation. From this experience, we have developed an optimization methodology based on an increased arithmetic intensity and applied it on acoustic and elastic wave equation for both isotropic and anisotropic media.

The comprehension of the subsurface structure is one of the important key aspect of the exploration process. Seismic depth imaging methods such as Kirchhoff, RTM and FWI are used to produce images of the subsurface allowing interpreter to make sense of the underlying structure and geology. RTM and FWI are heavily used for this task. Both are compute intensive applications and are the main consumer of high performance computing resources. It is therefore critical to choose the best hardware technology and optimize those applications to run efficiently in a minimum amount of time.

Heterogeneous architectures such as GPU have demonstrated during the past decade to be highly efficient and provide better performance than conventional CPU for seismic depth imaging. To be taken advantage of, GPUs requires the use of specific language extension that are considered as low-level programming model. CUDA is one of those and has been very popular for GPU programming. Unfortunately, CUDA is not standard and not portable on multiple vendors and hardware. Hence, OpenCL has been created to standardize the programming of accelerators and also aims to be portable on various hardware from CPU, GPU to FPGA. From a scientific user point of view, OpenCL is considered as CUDA as a low-level programming model. OpenACC is then born to simplify the programming of heterogeneous architecture. It is standard high-level programming model based on directives, similar to OpenMP. Despite offering a more productive way to develop parallel code on GPU, OpenACC has been considered way less efficient than low-level programming models like CUDA or OpenCL. In order to understand what is the performance gap between those models, we have implemented finite difference wave equation propagators using OpenACC. In this paper, we are demonstrating that OpenACC can achieve up to 70% of efficiency when compared to the best CUDA implementation. From this experience, we have developed an optimization methodology based on an increased arithmetic intensity and applied it on acoustic and elastic wave equation for both isotropic and anisotropic media. Moreover, those kernels are portable on a variety hardware of GPU: K40, K80 and P100, as well as CPU. In conclusion, OpenACC is a productive way to implement parallel computation and GPU, and languages are no longer a barrier to performance while algorithmic changes to increase arithmetic are key.