AAPG Datapages/Search and Discovery Article #90255 © 2016 AAPG/SEG International Conference & Exhibition, Barcelona, Spain, April 3-6, 2016

A New Analytical Model for Stencil-Based Seismic Algorithms Implementations on GPU

Dorfell Parra[1], William Salamanca[1], and Ana Ramirez[1]

[1]Universidad Industrial de Santander

Abstract

Seismic modeling is the basis for algorithms such as Reverse Time Migration (RTM). Among different seismic modeling strategies, the numerical method known as the Finite-Differences Time-Domain (FDTD) representation of the wave equation has been used to model the wave propagation phenomena.

Because the FDTD method is a highly computationally expensive method, the use of many-core architectures such as Graphical Processor Units (GPU) for large scale problems has become attractive. Previous works have achieved speedup up to 30x for seismic RTM naive implementations in comparison to CPU. The speed up in the execution time of an algorithm on a GPU depends on the implementation, the GPU hardware specifications, and the user preferences. Discerning which parameters highly affect the execution time will lead to better performance of the GPU implementation.

In order to find the best parameters configuration on the GPU, a traditional strategy is to start from a scratch implementation and run several tests to retrieve the bottlenecks of the implementation. Then, apply known optimization methods to avoid the bottlenecks and compare the performance. This strategy is however impractical for large scale problems such as the seismic RTM.

In this work, we propose an analytical model to estimate the performance, in terms of the execution time, of one stencil of the FDTD implementation. The stencil computation updates one spatial point on a grid by using only neighboring points on a regular spatial grid. The full solution of the FDTD method is then obtained by executing several stencils in parallel. The proposed analytical model uses GPU hardware and code implementation parameters, as well as micro-benchmarks to obtain information such as global memory access latency or number of floating point operations. The estimated execution time obtained with the proposed analytical model was compared with the measured execution time of the stencil implementation on an Nvidia Kepler K40 GPU. As future work this analytical model can be used in an optimization process of a RTM implementation.