Toward GPU Acceleration of the Reverse Time Migration Using Temporal Blocking

Long Qu[1], Hatem Ltaief[1], Hussain Salim[2], Thierry Tonellot[2], David Keyes[1]

[1]Kaust
[2]Aramco

Abstract

The oil and gas exploration industry uses the cutting-edge algorithm Reverse Time Migration (RTM) for seismic depth imaging in challenging geological settings. Using seismic datasets captured at various source and receiver locations, it solves the three-dimensional acoustic wave equation to provide high-resolution images. RTM's computational phases are predominantly composed of stencil computational kernels for the finite-difference time domain scheme, applying the absorbing boundary conditions, and I/O operations needed for the imaging condition. In this presentation, the full RTM workflow is integrated with the asynchronous Multicore Wavefront Diamond (MWD) tiling technique. By combining spatial with Temporal Blocking (TB) during the stencil computations, MWD offers significant increase in data reuse. The conventional synchronous RTM workflow faces new difficulties as a result of such integration since it requires rethinking of how the absorbing boundary conditions, the I/O operations, and the imaging condition operate. These disruptive changes are necessary to maintain the performance superiority of the asynchronous stencil execution during the time integration without deteriorating the quality of the subsurface image. By comparing the novel MWD-based RTM to the conventional Spatial Blocking (SB)-based RTM on a variety of shared-memory systems using syntactic datasets, we achieve up to 70% performance speedup. We further port and adapt our MWD-based RTM on massively parallel GPU architectures. Multiple challenges have been tackled. Besides the performance improvement of various compute kernels provided by CUDA, the existing MultiLayered Buffer Storage (MLBS) framework has been adapted and integrated into our implementation in order to asynchronously cache and prefetch partial snapshots data from/to various storage media. Such design transforms the original I/O-bound behavior of our GPU porting, and shifts it closer to a GPU compute-bound regime of execution. To our knowledge, this paper highlights for the first time the applicability of asynchronous executions with temporal blocking throughout the whole RTM on heterogeneous architectures. This may eventually create new research opportunities in improving the hydrocarbon extraction for the petroleum industry.