High-order Finite Difference Seismic Modeling on Reconfigurable Computing Platform

Chuan He, Department of Electrical Engineering, Guan Qin, Institute for scientific computation, and Wei Zhao, Department of Computer Science, Texas A&M University, College Station, TX 77843

Extended Abstract

Simulating linear wave propagation phenomenon numerically in 2D or 3D space with Finite Difference (FD) methods, although computationally intensive, are capable of dealing with more complex geological models than Fourier transformation methods or asymptotic ray-tracing methods. During the past decade, FD-based numerical seismic modelling efforts have grown rapidly in geophysics area with significant performance improvements in commodity computers and parallel computing environments. Pure software methods, from high-level parallelism on PC-Cluster system to low-level memory and disk optimization or even instruction-reordering are exhausted to accelerate the execution of these simulation tasks. However, they are still time-consuming procedures so cannot be used routinely except in institutes that can afford the high cost of running and maintaining supercomputers or large PC-cluster systems.

In this work, we proposed an FPGA-based solution to accelerate the execution of seismic acoustic and elastic modelling problems. The fundamental hindrance of this class of tasks is the massive data volume along with complex numerical algorithms. Specifically, memory bandwidth available between the computing engine (FPGA) and onboard memory modules has been proven a bottleneck preventing people taking full advantage of FPGA's computational potential. Migrating software version of the standard second-order FD method directly into RC platform, although may lead to impressive speeding-up comparing with contemporary commodity computers would eventually hit this memory bandwidth bottleneck and after that, no more accelerations could be achieved.

Our solution alleviates this performance bottleneck with two approaches: First, we adopt high-order spatial FD stencils and high-order time-integration schemes to increase the computational complexity of our numerical algorithm. Because the clock frequency applied to FD computing engine and external memory modules are within the same range at hundreds of million Hz, the bandwidth of onboard memories would be saturated rapidly with considerable RC resources inside FPGA being wasted. By choosing high-order FD schemes, we can always make the computations as complex as necessary to alter the performance bottleneck back to the computing engine. Moreover, high-order FD schemes allow larger sampling interval so that the total number of spatial grid points is considerably reduced. Consequently, the memory bandwidth requirement for the same problem decreases in an indirect way.

Second, an efficient on-chip sliding window buffering structure is constructed between the FD computing engine and onboard memory modules using RAM blocks inside FPGA chip. By exploring data dependency properties of our numerical algorithm, this buffering system is capable of providing a new set of input operands to the FD computing engine at every clock cycle. Although operand stencils of high-order schemes are much wider than the standard second-order method, the number of external memory accesses to evaluate one time-marching step at each grid point is kept unchanged. The only costs we pay for high-order schemes are on-chip memory blocks and conventional arithmetic units, which are all abundant inside an up-to-date high density FPGA chip.

In summary, by keeping the bandwidth of external memory modules always saturated, we can adjust the accuracy order of the FD scheme such that the utilization of reconfigurable resources on RC platform are also maximized. This advantage of our solution to this class of numerical algorithms makes it accommodated to a wide range of commercial RC platforms.