An implementation of the acoustic wave equation on FPGAs

MAXELER Technologies

This is a summary of a paper presented at the 78th Society of Exploration Geophysicists (SEG) Meeting, Las Vegas, USA, November 2008.

About the Authors

Tamas Nemeth, Joe Stefani and Wei Liu are Geophysicists at Chevron.
Oliver Pell is VP of Engineering and Rob Dimond is Chief Hardware Architect at Maxeler Technologies.
Ray Ergas is a Geophysicist, formerly at Chevron.

Summary

This study investigates the use of FPGAs as co-processors to accelerate an acoustic isotropic modeling application. Using a Maxeler MAX2 PCI Express x16 card based on FPGAs, we have achieved speedups of 2 *orders of magnitude* when compared to a single-core implementation on a modern CPU.

The acoustic modeling application in consideration is 3D finite difference, with 4th-order in time, 12th order in space, and uses single-precision floating-point arithmetic. Input data consist of two 3D earth model arrays (velocity and density) and a source function. The application iterates for a number of time steps with three wavefield arrays.

Total memory requirement on the CPU is therefore to store 4-bytes per point, for 5 arrays, or $20 * N^3$, where N is a spatial dimension. Since acceleration has the biggest impact for large projects, we assume that any solution should work for model sizes of $N = 1000$, with a total memory of 20 GB or more.

The acoustic variable density modeling code contains a kernel which consumes the majority of the compute cycles, indicating that the algorithm is a good candidate for acceleration.

We started the effort by studying the runtime performance of the application on CPUs. For a 300^3 mesh run on the AMD Opteron, the portion of runtime not contained within the kernel is only 0.04%. The kernel itself can be broken down into 5 subcomponents, which can be considered separately for acceleration.

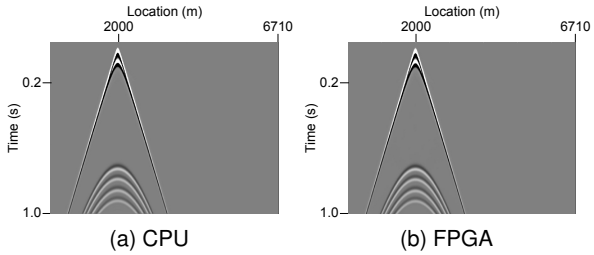

(a) CPU (b) FPGA

Figure 1: Comparison of the modeled seismograms between the CPU and the FPGA implementations.

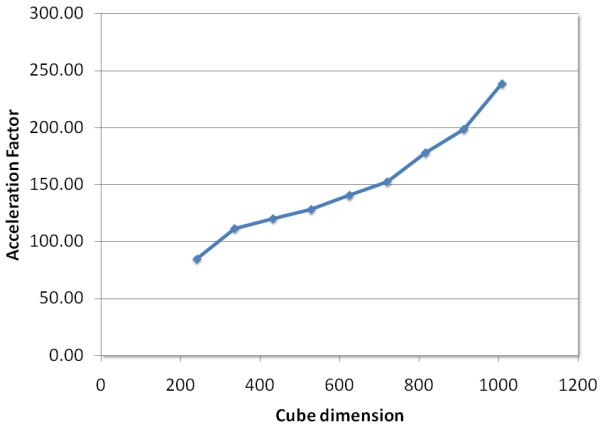

Figure 2: Acceleration over a Single Core CPU Implementation.

Figure 1 shows the results of testing a 2-layer, 3D model, with $672*672*256$ grid size. Seismograms were calculated both on CPU using the reference modeling code (a) and on the FPGA boards using the FPGA implementation (b).

The velocities of the 2 layers were 1500 and 4000 m/s and the density contrast was 1:4. The source was placed at location (2000m, 500m) and the receiver line was placed at depth of 550 m. The modeling parameters were a grid size of 10 m, peak frequency of 32 Hz and 643 time steps.

Figure 2 shows a performance comparison for several different sizes of cube between the FPGA implementation and the reference software running on a single core of an Intel Xeon clocked at 2.66GHz. The FPGA provides a speedup of nearly 250x over a single core for a large cube.

The implications of these results are twofold:

- Non-standard HPC technologies, such as FPGA chips can play a useful role in seismic data processing.
- The traditional notion of what is computationally cheap or expensive can be severely altered with the emerging new hardware and software technologies.