Acceleration of the Anisotropic PSPI Imaging Algorithm with DFEs

MAXELER Technologies

This is a summary of a paper presented at the 82nd Society of Exploration Geophysicists (SEG) Meeting, Las Vegas, September 2012.

About the Authors

Oliver Pell is VP of Engineering and Diego Oriato is a Distinguished Acceleration Architect at Maxeler Technologies. Daniela Theis, Guido Satta and Ernesto Bonomi are with CRS4 Imaging and Numerical Geophysics, Italy. Carlo Tomas and Luca Cazzola are with ENI E&P Division,Italy.

Summary

Seismic computation benefits from new hardware platforms through faster application and implementation of more complex and accurate models. However, the difficulty lies in restructuring the algorithm to efficiently utilize the new hardware. Frequency-domain depth migration uses two-dimensional Fourier transform as a mathematical tool. After Fourier transform, the data are decomposed into a series of frequency panels that are processed independently in a dataflow engine (DFE). The first key factor to achieve optimal performance is to utilize all the code parallelism available in the PSPI algorithm. The second step is to apply bespoke wavefield compression.

Seismic imaging in a stratified medium can be done with plane Phase-Shift, but the case with lateral velocity variations requires more effort. The propagation model must be modified in order to construct a pure spectral method for downward continuation in an inhomogeneous isotropic medium.

A DFE can be regarded as programmable hardware logic connected to a local memory. The separation of computation and communication allows for full usage of available hardware resources. We determined all six PSPI components had to be processed using the DFE. Due to the fact that each instruction occupies physical space on a DFE, the accelerated implementation used a maximum number of reference velocities that could be computed on a single stream of data. Computation of the last reference velocity was performed twice to ensure that the interpolation block always had two consecutive reference velocities available.

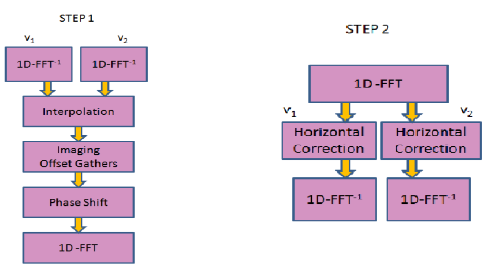

(a) I/O-bound (b) Compute-bound

Figure 1: Workflow for PSPI with two reference velocities

The 2D-FFT step was executed with 1D-FFT hardware logic blocks. The most efficient way to perform data transposition is to write data to the DFEs memory in one direction and read it back in the opposite one. This resulted in a single depth migration step implemented by two streaming steps, shown in Figure 1. The complex source and receiver wavefields were transferred from the CPU to the DFE memory, streamed into the hardware logic, extrapolated in parallel and combined together in the imaging logic.

Two generations of Maxeler hardware were employed: the MaxNode-1842 utilizing 4 MAX2 DFEs and the MaxNode-1843 with 4 MAX3 DFEs. In order to evaluate performance, the MAX2 and MAX3 systems are compared to the original MPI/Fortran version. The speedup increases with the number of reference velocities of the layer, observable in Figure 2. This constitutes an important factor for the migration with DFEs of real seismic data.

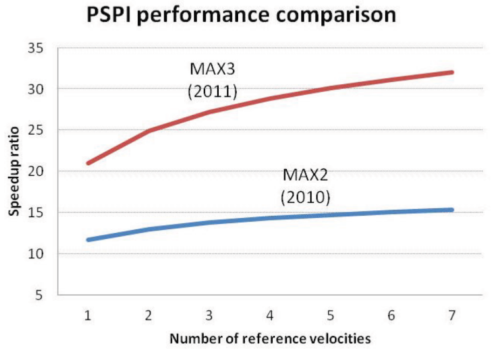

Figure 2: MAX2 and MAX3 nodes compared to 8 Intel Xeon "Westmere" 2.6GHz cores

Original Publication Date: September 2012