

FD modeling beyond 70Hz with FPGA acceleration

This is a summary of a talk presented at the SEG 2010 HPC Workshop in Denver, USA, October 2010.

About the Authors

Oliver Pell is VP of Engineering at Maxeler Technologies. Diego Oriato is an acceleration architect at Maxeler Technologies.

Clara Andreoletti and Nicola Bienati are scientists at ENI S.p.A. Exploration and Production Division, AESI Department.

Summary

One of the main factors that determine the resolution of seismic images is the bandwidth of the seismic wavelet. Finite difference (FD) modeling and Reverse Time Migration (RTM) encounter particular problems increasing the wavelet bandwidth at the upper end of the spectrum because of the large impact this has on the computation resource requirements. Increasing the upper modeled frequency requires a finer spatial sampling while the CFL limit implies that the modeling time step must decrease. The net result is that the number of arithmetic operations grows as the fourth power of the frequency (shown in [Figure 1](#)).

We use hardware acceleration to increase the maximum modeled frequency without a significant increase in compute time. The Maxeler MAX2 FPGA card used has two Xilinx Virtex-5 FPGAs and 24GB of on-board DRAM. MAX2 has a high-speed MaxRing link between cards that enables multiple MAX2 cards to work together to achieve the 70Hz bandwidth objective.

A particular bandwidth and number of grid points, N , translate into a memory requirement depend-

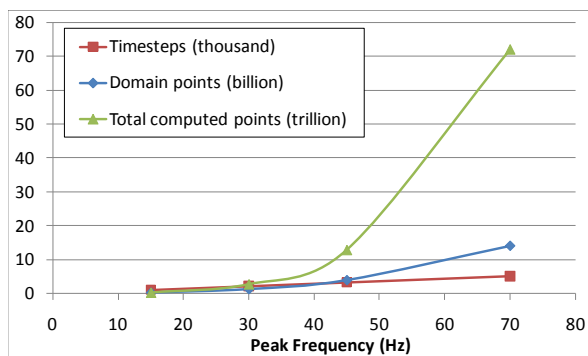


Figure 1: Computation and memory cost of FD modeling at high frequencies.

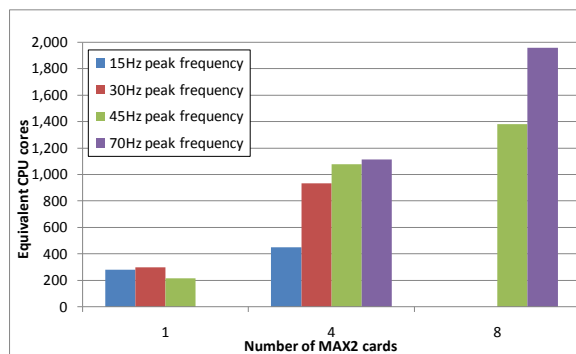


Figure 2: Performance of the FPGA-accelerated modeling.

ing on the specific forward modeling propagator and RTM algorithm. We use an acoustic isotropic propagator, which requires memory for one model volume and two pressure volumes (for the current and previous timesteps), giving a memory requirement for forward modeling of $3 \times N \times BytesPerItem$. To support a 70Hz peak frequency for a representative 17.5km x 20.3km x 7.8km shot aperture requiring 14 billion grid points, we need between 4 and 8 MAX2 cards to hold the volume, depending on the RTM scheme used. The application uses all the 24GB of memory on each MAX2 card and also compresses the stored data to further reduce the total data volume.

We use one standard x86 server connected to two 1U MaxBoxes, each containing 4 MAX2 cards. This provides 16 FPGAs and 192GB of FPGA card memory in a dense 3U form factor. For processing, we decompose the problem domain in one dimension between the FPGAs. Boundary (halo) regions are communicated between the FPGAs each timestep, and the one dimensional decomposition means that FPGAs only need to send/receive data from their neighbors to the left and right, which is implemented efficiently using the MaxRing interconnect. Communication is overlapped with the computation, so there is no significant performance impact.

We compare the performance of the MAX2 system to a C++ software version on a cluster with 32 3GHz X86 cores communicating via MPI over Infiniband ([Figure 2](#)). The maximum performance of the accelerated node is equivalent to nearly 2000 CPU cores: one MAX2 card provides the equivalent performance of over 200 CPU cores.