Geophysical Research Abstracts
Vol. 19, EGU2017-12381, 2017
EGU General Assembly 2017

Fast in-memory elastic full-waveform inversion using consumer-grade GPUs

Tore Sivertsen Bergslid, Espen Birger Raknes, and Børge Arntsen

Department of Geoscience and Petroleum, Norwegian University of Science and Technology, Trondheim, Norway (tore.bergslid@ntnu.no)

Full-waveform inversion (FWI) is a technique to estimate subsurface properties by using the recorded waveform produced by a seismic source and applying inverse theory. This is done through an iterative optimization procedure, where each iteration requires solving the wave equation many times, then trying to minimize the difference between the modeled and the measured seismic data. Having to model many of these seismic sources per iteration means that this is a highly computationally demanding procedure, which usually involves writing a lot of data to disk. We have written code that does forward modeling and inversion entirely in memory.

A typical HPC cluster has many more CPUs than GPUs. Since FWI involves modeling many seismic sources per iteration, the obvious approach is to parallelize the code on a source-by-source basis, where each core of the CPU performs one modeling, and do all modelings simultaneously. With this approach, the GPU is already at a major disadvantage in pure numbers. Fortunately, GPUs can more than make up for this hardware disadvantage by performing each modeling much faster than a CPU.

Another benefit of parallelizing each individual modeling is that it lets each modeling use a lot more RAM. If one node has 128 GB of RAM and 20 CPU cores, each modeling can use only 6.4 GB RAM if one is running the node at full capacity with source-by-source parallelization on the CPU. A parallelized per-source code using GPUs can use 64 GB RAM per modeling. Whenever a modeling uses more RAM than is available and has to start using regular disk space the runtime increases dramatically, due to slow file I/O.

The extremely high computational speed of the GPUs combined with the large amount of RAM available for each modeling lets us do high frequency FWI for fairly large models very quickly. For a single modeling, our GPU code outperforms the single-threaded CPU-code by a factor of about 75. Successful inversions have been run on data with frequencies up to 40 Hz for a model of 2001 by 600 grid points with 5 m grid spacing and 5000 time steps, in less than 2.5 minutes per source. In practice, using 15 nodes (30 GPUs) to model 101 sources, each iteration took approximately 9 minutes. For reference, the same inversion run with our CPU code uses two hours per iteration. This was done using only a very simple wavefield interpolation technique, saving every second timestep. Using a more sophisticated checkpointing or wavefield reconstruction method would allow us to increase this model size significantly.

Our results show that ordinary gaming GPUs are a viable alternative to the expensive professional GPUs often used today, when performing large scale modeling and inversion in geophysics.