Geophysical Research Abstracts
Vol. 17, EGU2015-9993, 2015
EGU General Assembly 2015

Wavefield Compression for Full-Waveform Inversion

Christian Boehm (1), Andreas Fichtner (1), Josep de la Puente (2), and Mauricio Hanzich (2)

(1) Institute of Geophysics, ETH Zurich, Zurich, Switzerland (christian.boehm@erdw.ethz.ch), (2) Barcelona Supercomputing Center, Barcelona, Spain

We present compression techniques tailored to iterative nonlinear minimization methods that significantly reduce the memory requirements to store the forward wavefield for the computation of sensitivity kernels.

Full-waveform inversion on 3d data sets requires massive computing and memory capabilities. Adjoint techniques offer a powerful tool to compute the first and second derivatives. However, due to the asynchronous nature of forward and adjoint simulations, a severe bottleneck is introduced by the necessity to access both wavefields simultaneously when computing sensitivity kernels. There exist two opposing strategies to deal with this challenge. On the one hand, conventional approaches save the whole forward wavefield to the disk, which yields a significant I/O overhead and might require several terabytes of storage capacity per seismic event. On the other hand, checkpointing techniques allow to trade an almost arbitrary amount of memory requirements for a – potentially large – number of additional forward simulations.

We propose an alternative approach that strikes a balance between memory requirements and the need for additional computations. Here, we aim at compressing the forward wavefield in such a way that (1) the I/O overhead is reduced substantially without the need for additional simulations, (2) the costs for compressing/decompressing the wavefield are negligible, and (3) the approximate derivatives resulting from the compressed forward wavefield do not affect the rate of convergence of a Newton-type minimization method. To this end, we apply an adaptive re-quantization of the displacement field that uses dynamically adjusted floating-point accuracies – i.e. a locally varying number of bits – to store the data. Furthermore, the spectral element functions are adaptively downsampled to a lower polynomial degree. In addition, a sliding-window cubic spline re-interpolates the temporal snapshots to recover a smooth signal. Moreover, a preprocessing step identifies spatial and temporal “shadow zones” where storing the forward wavefield is not required at all, i.e. regions with a negligible magnitude of the gradient as well as time steps prior to the first arrival of the forward or adjoint wavefront.

It is important to note that a poorly approximated wavefield might result in an approximation of the gradient that is not a direction of ascent. To overcome this difficulty, we provide error estimates and outline a strategy to adaptively control the rate of compression during the iterations. We present numerical results that compare the performance and quality of the compression.