11th World Congress on Computational Mechanics (WCCM XI)
5th European Conference on Computational Mechanics (ECCM V)
6th European Conference on Computational Fluid Dynamics (ECFD VI)
July 20–25, 2014, Barcelona, Spain

SCALING SEISMIC IMAGING ALGORITHMS TO PETASCALE COMPUTING AND BEYOND

Matthieu Lefebvre[1], Ebru Bozdağ[2], Henri Calandra[3], Dimitri Komatitsch[4], Wenjei Lei[5], Daniel Peter[6], Herurisa Rusmanugroho[7], James Smith[8] and Jeroen Tromp[9]

[1] Princeton University, ml15@princeton.edu
[2] Université de Nice Sophia Antipolis, bozdag@princeton.edu
[3] Total E&P, henri.calandra@total.com
[4] CNRS, komatitsch@lma.cnrs-mrs.fr
[5] Princeton University, lei@princeton.edu
[6] ETH Zürich peterda@ethz.ch
[7] Princeton University, herurisa@princeton.edu
[8] Princeton University, jas11@princeton.edu
[9] Princeton University, jtromp@princeton.edu

Key words: *Geophysics, Applications, Computing Methods, GPU computing, Workflow Management, Parallel I/Os*

Knowledge about Earth's interior comes mainly from seismic observations and measurements. Seismic tomography is the most powerful technique for determining 3D images of the Earth —usually in terms of wavespeeds, density, or attenuation— using seismic waves generated by earthquakes or man-made sources recorded by a set of receivers. Advances in the theory of wave propagation and 3D numerical solvers together with dramatic increases in the amount and quality of seismic data and rapid developments in high-performance computing offer new opportunities to improve our understanding of the physics and chemistry of Earth's interior. Adjoint methods provide an efficient way of incorporating 3D numerical wave simulations in seismic imaging, and have been successfully applied for regional- and continental-scale problems [1] and —to some extent— in exploration seismology [2]. However, it has so far remained a challenge on the global scale and in 3D exploration, mainly due to computational limitations.

In the context of adjoint tomography, scientific workflows are well defined. They consist of a few collective steps (e.g., mesh generation, model updates, etc.) and of a large number of independent steps (e.g., forward and adjoint simulations for each seismic event, pre- and post-processing of seismic data, etc.). The goal is to increase the accuracy of seismic models while keeping the entire procedure as efficient and stable as possible.

The forward and adjoint simulations have been optimized to perform on multi- and manycores architectures, in particular GPUs [3]. While computational power still remains an important concern, large-scale experiments and big data sets create bottlenecks in workflows causing significant I/O problems on HPC systems.

Legacy seismic data formats were initially designed for specific seismic applications involving limited data sets, with little concern for performance. We are developing a new modern seismic data format based on ORNL's ADIOS libraries –called the Adaptable Seismic Data Format (ASDF)– that is suited for a variety of seismic workflows, allowing users to retain provenance related to observed and simulated seismograms. The preprocessing tools (resampling, filtering, window selection, computing adjoint sources, etc.) are modified and parallelized to take advantage of this new data format. We accommodate the ADIOS libraries in our numerical solvers to reduce the impact of accessing files on disk during simulations (i.e., meshes, kernels, models, etc.) and also during other processing steps. We adjust post-processing tools (i.e., summing, pre-conditioning and smoothing gradients, model updates, etc.) accordingly. Moreover, parallel visualization tools, such as VisIt, take advantage of metadata included in our ADIOS outputs to extract features and display massive datasets. We also investigate the possibility of automatizing the whole workflow with the help of a scientific management workflow system such as Swift [4]. With our improvements on seismic data formats and file I/O during the inversion workflow, researchers will be able to conduct large-scale seismic inversions much more effectively on petascale HPC systems. However, further improvements on large-scale seismic inversion workflows on HPC systems are still needed.

REFERENCES

[1] H. Zhu, E. Bozdağ, D. Peter, and J. Tromp, "Structure of the European upper mantle revealed by adjoint tomography," *Nature Geoscience*, vol. 5, no. 7, pp. 493–498, 2012. [Online]. Available: `http://www.nature.com/doifinder/10.1038/ngeo1501`

[2] Y. Luo, J. Tromp, B. Denel, and H. Calandra, "3D coupled acoustic-elastic migration with topography and bathymetry based on spectral-element and adjoint methods," *GEOPHYSICS*, vol. 78, no. 4, pp. S193–S202, 2013. [Online]. Available: `http://library.seg.org/doi/abs/10.1190/geo2012-0462.1`

[3] M. Rietmann, P. Messmer, T. Nissen-Meyer, D. Peter, P. Basini, D. Komatitsch, O. Schenk, J. Tromp, L. Boschi, and D. Giardini, "Forward and adjoint simulations of seismic wave propagation on emerging large-scale GPU architectures," in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, ser. SC '12. Los Alamitos, CA, USA: IEEE Computer Society Press, 2012, pp. 38:1—-38:11. [Online]. Available: `http://dl.acm.org/citation.cfm?id=2388996.2389048`

[4] M. Wilde, I. Foster, K. Iskra, P. Beckman, Z. Zhang, A. Espinosa, M. Hategan, B. Clifford, and I. Raicu. Parallel Scripting for Applications at the Petascale and Beyond. *Computer*, 42(11):50–60, 2009.