Pillars of the Mantle: Imaging the Interior of the Earth with Adjoint Tomography

David Pugmire
Oak Ridge National Laboratory
pugmire@ornl.gov

Ebru Bozdağ
Colorado School of Mines

Matthieu Lefebvre
Princeton University

Jeroen Tromp
Princeton University

Dmitri Komatitsch
University of Aix-Marseille

Daniel Peter
King Abdullah University of Science
and Technology

Norbert Podhorszki
Oak Ridge National Laboratory

Judith Hill
Oak Ridge National Laboratory

ABSTRACT

In this work, we investigate global seismic tomographic models obtained by spectral-element simulations of seismic wave propagation and adjoint methods. Global crustal and mantle models are obtained based on an iterative conjugate-gradient type of optimization scheme. Forward and adjoint seismic wave propagation simulations, which result in synthetic seismic data to make measurements and data sensitivity kernels to compute gradient for model updates, respectively, are performed by the SPECFEM3D_GLOBE package [1] [2] at the Oak Ridge Leadership Computing Facility (OLCF) to study the structure of the Earth at unprecedented levels. Using advances in solver techniques that run on the GPUs on Titan at the OLCF, scientists are able to perform large-scale seismic inverse modeling and imaging. Using seismic data from global and regional networks from global CMT earthquakes, scientists are using SPECFEM3D_GLOBE to understand the structure of the mantle layer of the Earth. Visualization of the generated data sets provide an effective way to understand the computed wave perturbations which define the structure of mantle in the Earth.

ACM Reference format:

David Pugmire, Ebru Bozdağ, Matthieu Lefebvre, Jeroen Tromp, Dmitri Komatitsch, Daniel Peter, Norbert Podhorszki, and Judith Hill. 2017. Pillars of the Mantle: Imaging the Interior of the Earth with Adjoint Tomography. In *Proceedings of PEARC17, New Orleans, LA, USA, July 09-13, 2017*, 4 pages. <https://doi.org/10.1145/3093338.3104170>

1 INTRODUCTION

Knowledge about the interior of the Earth 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 the interior of the Earth. 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 [3] and, to some extent, in exploration seismology [5]. However, it has so far remained a challenge on the global scale and in 3D exploration, mainly due to computational limitations.

2 COMPUTATIONAL SEISMOLOGY

The steady increase in data quality and the number of global seismographic stations have substantially grown the amount of data available for construction of Earth models. Meanwhile, developments in the theory of wave propagation, numerical methods and HPC systems have enabled unprecedented simulations of seismic wave propagation in realistic 3D Earth models which lead the extraction of more information from data, ultimately culminating in the use of entire three-component seismograms.

A major challenge in computational seismology is to image the entire planet using adjoint tomography. Due to the vast amount of high-quality seismic data that can potentially be assimilated, and the intense computational requirements, global imaging is one of the extreme challenges in seismology.

Initial work has been performed using low resolution global inversions using 253 earthquakes for a transversely isotropic crust and mantle model on the Oak Ridge Leadership Computing Facility's Cray XK7 Titan system. Recent improvements in 3D solvers, such as the GPU version of the SPECFEM3D_GLOBE package, has allowed higher-resolution, and higher frequency body waves to be used in the simulations.

Recently generated results [4] already indicate several prominent features reported in high-resolution continental studies, such as major slabs (Hellenic, Japan, Bismarck, Sandwich, etc.) and enhancement in plume structures (the Pacific superplume, the Hawaii hot spot, etc.). The ultimate goal is to assimilate seismic data from more than 6,000 earthquakes within the magnitude range between 5.5 and 7.0.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

PEARC17, July 09-13, 2017, New Orleans, LA, USA

© 2017 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5272-7/17/07.

<https://doi.org/10.1145/3093338.3104170>

3 SIMULATION AND DATA MANAGEMENT

We have started global inversions with a carefully selected set of 253 earthquakes, which means we need to perform 253 forward and adjoint simulations to compute the gradient for each iterations. We have so far performed 10 low-resolution (minimum period in simulations $> 27s$) global iterations with 100 min-length seismograms on Titan using the GPU version of SPECfEM3D_GLOBE package. For 100 min low-resolution simulations, 253 jobs for each event are run through a batch submission on 12,650 GPU where forward and adjoint simulations take about 10 and 30 min, respectively. Our aim is to assimilate all earthquakes in the catalogue which is around 6,000 earthquakes since 1999 within the magnitude range of 5.5 to 7.0 and go down to 9 s resolution in inversions with longer duration simulations (200 m) to take the advantage of high-frequency body waves and major-arc seismic waves, respectively.

To take full advantage of these simulations on the resources at the Oak Ridge Leadership Computing Facility, a solid framework is required for managing big data sets during pre-processing (e.g., data requests and quality checks), gradient calculations, post-processing (e.g., pre-conditioning and smoothing gradients), and visualization.

One of the bottlenecks in this workflow was the tremendous number of files that had to be read, and written during the simulations. Each earthquake requires a very large number of individual seismogram data files. To alleviate these bottlenecks, a new file format (ASDF, Adaptable Seismic Data Format [9]) was developed which stores all data relevant to an earthquake in a single data file. This new format was developed on top of the Adaptable I/O System (ADIOS) [8]. ADIOS provides a scalable I/O method for handling such datasets efficiently in parallel at any scale. As a result of this effort, the reading and writing of earthquake data is no longer a bottleneck in the tomography workflow.

4 VISUALIZATION

The simulations were rendered on the post-processing cluster Rhea at the OLCF using two different software packages: VisIt [6] and Blender [10]

First, VisIt was used for the initial exploration of the data and for extraction of the seismological features. The broad capability of VisIt, and it's demonstrated scalability [7], proved valuable in the ability to experiment with different visualiation techniques, and in the creation of timely results. Utilizing VisIt's plugin-architecture, a data reader plugin was developed which reads the ADIOS format simulation output files. Once developed, the full functionality of VisIt can be used for analysis and visualization of the data.

Second, Blender was used for the setup of lighting, materials, camera paths and rendering of the the geometry. Python scripting was used to control the orchestration of different geometry, as well as camera animation.

This work involved the exploration of the variable $dvsv$, which is the shear-wave perturbations with respect to the 1D reference model. Postive and negative perturbations indicate shear-wavespeeds faster and slower than the average share-wavespeed of the mantle given by the 1D refernece model.

Figure 1: Whole earth volume rendering of shear-wave perturbations.

The computational mesh consists of approximately 100 million unstructured zones, and decomposed into several hundred spatial domains. The variables are defined at the nodes.

4.1 Workflow Integration

VisIt supports a plug-in based architecture which makes it straightforward to develop for arbitrary data formats. In order to support the SPECfEM3D_GLOBE workflow, a parallel data reading plugin for VisIt was developed that reads the ADIOS format data. Once read in by the plugin, the large number of operations inside VisIt can be used for visualization and analysis.

4.2 Data Visualization

Figure 1 shows a volume rendering of the shear-wave perturbations variable in the 3D coordinate system which is used during simulations. Continental boundaries (shown in white), and seismic hot-spots (showin in red) are rendered to provide spatial context. The transfer function for the shear-wave perturbations are set to highlight values between $\pm 4\%$, which can highlight seismological features of interest.

In order to better see the internal structure within the Earth, a change of coordinate system operation was performed, and shown in Figures 2. The X,Y,Z cartesian coordinates of the mesh were transformed to spherical coordinates, R, θ, ϕ , and then mapped back to cartesian coordinates with R as the height, and θ and ϕ as the vertical and horizontal directions. As before, continental boundaries, and seismic hot-spots are shown to provide spatial context. Movie frames were generated by animating the shear-wave perturbations values between $\pm 8\%$.

Because the radial dimension is now visible, it is much easier to see the structure inside the Earth. This particular set of visualizations is especially valuable to seismologists as it allows them to better observe the features and structure of the Earth, how they

Figure 2: Iso contour rendering of positive and negative shear-wave perturbations with a spherical coordinate transform.

Figure 3: Rendering of land and water features along with seismological features. The land and water features are lowered down over time to reveal the seismological features.

develop radially, how they interact with neighboring features, and where they are spatially.

4.3 Video Script

The video begins with a sequence to orient the scientist, as shown in Figure 3. Following this orientational sequence, the labeling of significant hotspots along with a semi transparent plane at 660km are displayed in Figure 4. These labels allow the scientists to see the relationships of the seismological features to important seismic hotspots.

The final sequence is a fly through of the seismological features which allows scientists to better see and analyze the plume structures in the simulation. Of particular importance to seismologist is the Pacific Superplume, which is shown in Figure 5.

5 CONCLUSIONS AND FUTURE WORK

Department of Energy petascale resources at the Oak Ridge Leadership Computing Facility provide the capability to perform computational tomography of the Earth at unprecedented levels. Advanced

Figure 4: Rendering of the seismological features long with labels for seismic hotspots. The semi-transparent plane shows the 660km depth plane.

Figure 5: Rendering of a closeup view of the Pacific Superplume.

solvers, running on the GPUs in Titan, allows researches to use large-scale seismic inverse modeling and imaging at scales greater than previously possible, allowing researchers opportunities to better understand the structure of the mantle in the Earth.

Critical in this scientific workflow is the visualization of simulation results. The visualizations in this collaboration have been used to help verify that the code is running correctly, as well as explore the inherently three-dimensional, and complex structure of the interior of the Earth. In particular, the iso-contour, and horizontal slice renderings are a critical step to understanding the 3D geometry of structures such as the major plumes beneath Samoa and Africa, which originate from core-mantle boundary coming to the surface. After 15 iterations of the model, we have already started observing and commenting on long-debated subjects of Earth sciences.

Visualization has also been used to perform comparisons between different model iterations of the mantle of the Earth. In exploring the results of a particular model, a variety of techniques are typically used and highlighted in this video. These have been used to gain insight into existing data, and further techniques are

under development to gain further insight into the data being generated. Insights gained from these scientific visualizations have been used as catalysts for further computations, and analyses.

In the future we plan to automate many of the feature extraction, and rendering tasks and perform them as in situ tasks. Our plan is to implement these tasks as analysis and visualization services [?] using VTK-m [?] in the ADIOS middleware system.

6 ACKNOWLEDGEMENTS

This work was done using resources generously made available by the Oak Ridge Leadership Computing Facility located at the Oak Ridge National Laboratory.

REFERENCES

- [1] D. Komatitsch and J. Tromp. Spectral-element simulations of global seismic wave propagationI. Validation. *Geophys. J. Int.*, 149:390–412, 2002.
- [2] D. Komatitsch and J. Tromp. Spectral-element simulations of global seismic wave propagationII. Three-dimensional models, oceans, rotation and self-gravitation. *Geophys. J. Int.*, 150:308–318, 2002.
- [3] E. Bozdag, M. Lefebvre, W. Lei, D. Peter, J. Smith, D. Komatitsch, J. Tromp. Big Data and High-Performance Computing in Global Seismology. *Geophysical Research Abstracts*, Vol. 16, EGU2014-16606-1, 2014.
- [4] E. Bozdag, D. Peter, M. Lefebvre, D. Komatitsch, J. Tromp, J. Hill, N. Podhorszki, D. Pugmire (2016), Global Adjoint Tomography: First-Generation Model. *Geophysical Journal International*. 207(3): 1739-1766.
- [5] M. Lefebvre, E. Bozdag, H. Calandra, D. Komatitsch, W. Lei, D. Peter, H. Rusmanugroho, J. Smith and J. Tromp. Scaling Seismic Imaging Algorithms to Petascale Computing and Beyond. 11th World Congress on Computational Mechanics, July 20-25, 2014. Barcelona Spain.
- [6] H. Childs, E. Brugger, K. Bonnell, J. Meredith, M. Miller, B. Whitlock, N. Max. A Contract-Based System for Large Data Visualization. *IEEE Visualization 2005*, October 23-25 2005, Minneapolis, Minnesota.
- [7] H. Childs, D. Pugmire, S. Ahern, B. Whitlock, M. Howison, Prabhat, G. Weber, and E. W. Bethel. Extreme Scaling of Production Visualization Software on Diverse Architectures. *Computer Graphics and Applications*, special issue on Ultrascale Visualization, volume 30, number 3, pp. 22-31, May/June 2010.
- [8] Q. Liu, J. Logan, Y. Tian, H. Abbasi, N. Podhorszki, J. Cho, S. Klasky, R. Tchoua, J. Lofstead, R. Oldfield, M. Parashar, N. Samatova, K. Schwan, A. Shoshani, M. Wolf, K. Wu, W. Yu. (2014), Hello ADIOS: the challenges and lessons of developing leadership class I/O frameworks. *Concurrency Computat.: Pract. Exper.*, 26: 1453–1473.
- [9] <http://sdf.readthedocs.org>
- [10] <https://www.blender.org>
- [11] K. Moreland, C. Sewell, W. Usher, L. Lo, J. Meredith, D. Pugmire, J. Kress, H. Schroots, K. Ma, H. Childs, M. Larsen, C. Chen, R. Maynard, B. Geveci. VTK-m: Accelerating the Visualization Toolkit for Massively Threaded Architectures. *IEEE Computer Graphics and Applications*, 36(3):48-58, 2016.
- [12] D. Pugmire, J. Kress, J. Meredith, N. Podhorszki, J. Choi, S. Klasky. Towards scalable visualization plugins for data staging workflows. *Big Data Analytics: Challenges and Opportunities (BDAC-14) Workshop at Supercomputing Conference*, 2014.