H048

The EAGE 3D Anisotropic Elastic Modelling Project

W.A. Mulder* (Delft University of Technology), L. Nicoletis (Institut Français du Pétrole) & T.A. Alkhalifah (King Abdulaziz City for Science and Technology)

SUMMARY

Three-dimensional anisotropic elastic data computed by a finite-difference code on a two-dimensional section of the SEG/EAGE Overthrust model will enable benchmarking of processing and imaging algorithms and other three-dimensional modelling tools. The use of a two-dimensional model reduces the number of shots and the data volume. We present a description of the model, software and hardware, acquisition layout, and some initial results obtained on the IBM Blue Gene computer in Groningen.

Introduction

The project's goal is to provide three-dimensional synthetic data for an elastic anisotropic model. Realistic finite-difference simulations require PetaFlops hardware. One PetaFlops means 10^{15} floating point computations per second. This kind of performance is not available, although the first PetaFlops machine may appear this or next year. Because the computational cost of an explicit time-stepping scheme scales with fourth power of the maximum frequency in the data, one can always drop some realism to make the problem fit the available hardware. Memory requirements scale with the third power.

For this project, we have constructed a simplified two-dimensional model. Performing three-dimensional computations in a two-dimensional model reduces the number of required shots and keeps the data volume to a size that can be downloaded over a fast network.

The synthetic data will enable testing and development of processing and imaging procedures. In addition, three-dimensional modelling tools, as for instance based on ray tracing, can be benchmarked.

Model definition

A synthetic elastic survey requires a model of the subsurface in terms of density and P- and S-wave velocities and anisotropy parameters. A large number of shots must be fired for proper imaging. Each shot represents a separate computation. To reduce the number of shots and the data volume, three-dimensional computations are performed in a two-dimensional model.

The model was built from a two-dimensional section of the SEG/EAGE three-dimensional Overthrust model [1], starting from an initial index grid with a spacing of 1 meter. From this, models were built on coarser grids. Each index was assigned a density, P- and S-velocity, Thomsen parameters for transversely isotropic media, and a direction for the symmetry axis. The maximum P phase-velocity and the density are shown in Fig. 1. The total length is 20 km and the maximum depth 4 km, but only the left part of 10 km length was used in the numerical simulations.

Computing platform

The computations are performed on an IBM Blue Gene/L computer called "Stella" in Groningen, in the north of the Netherlands. This machine will be used for LOFAR, a project initiated by astronomers to build the world's largest radio telescope [4]. fast network will connect a large low-frequency array of small antennas. Radio interferometry will be carried out in real time. Geophysical and agricultural applications will also use the network. During

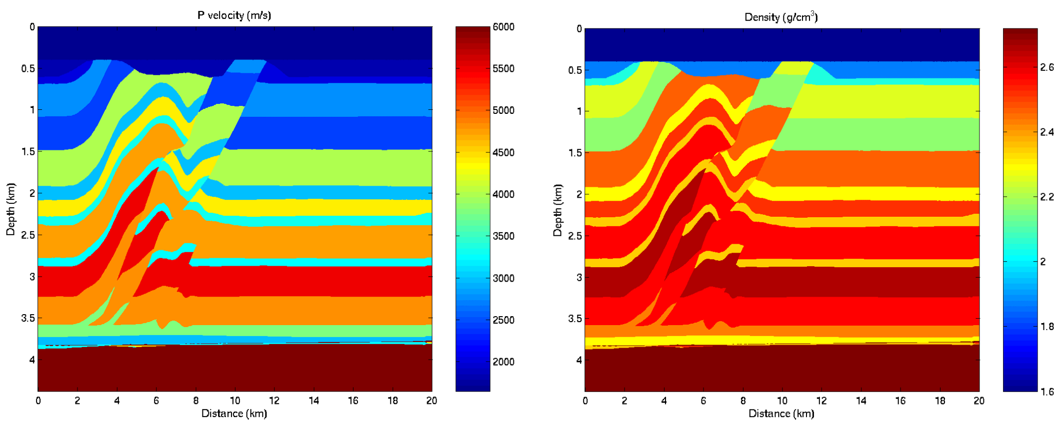

Fig. 1. Maximum P phase-velocity (left) and density (right). Only the 10 km part on the left was used in the simulations.

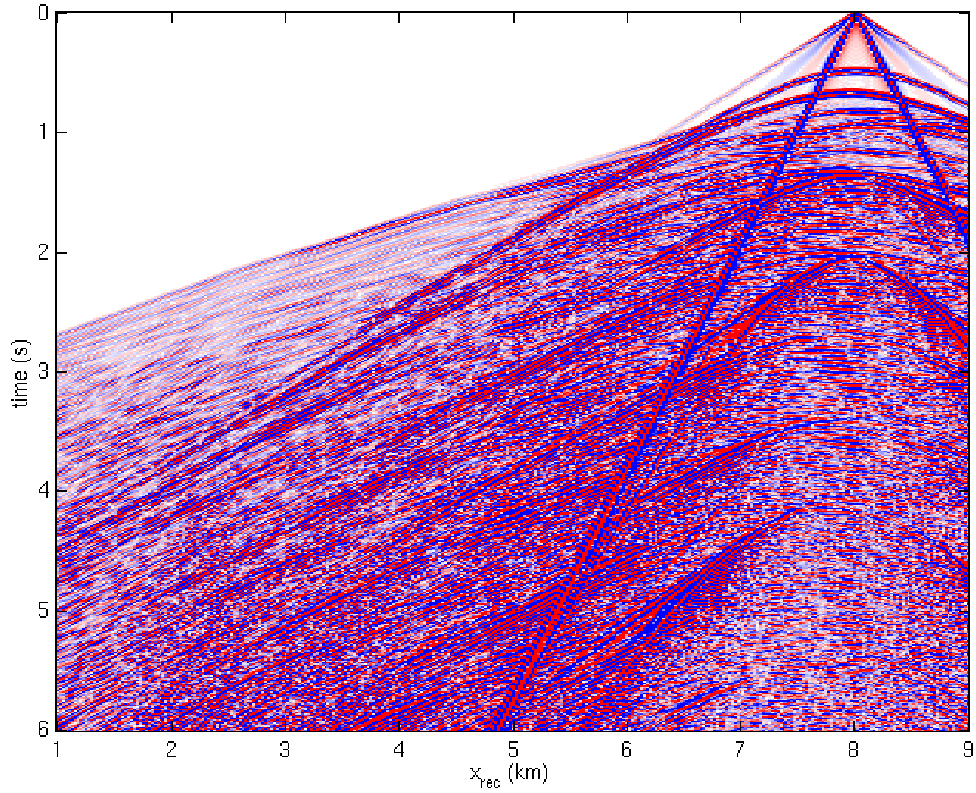

Fig. 2. Recorded vertical displacement. The data have been multiplied by time and are strongly clipped, to 1% of the maximum absolute value after time weighting.

the construction of the network and telescope, the computer is available for other applications.

The machine consists of six racks with 32 node boards, 16 compute cards per board, two chips and 512 Mbytes of memory per card. Each chip contains two CPUs and four FPUs (floating point units), leading to a total of 2×6144 CPUs. Each CPU runs at 700 MHz. With 2×4 FPUs, the theoretical peak performance is 8 Flops (floating point operations) per cycle per chip. In practice, one CPU is used for I/O. The other, together with its 4 FPUs, is used for the computations, so we only have half this performance, with a theoretical maximum of 34.4 TFlops for the full machine. The individual nodes have small memory and slow processing speeds (at most 2.8 GFlops) according to present-day standards. Their large numbers and fast internal network amply compensate for this. The network connects nearest neighbours on a three-dimensional 64×32×32 grid of nodes, with wrap-around at the ends. On top of that, there is a global tree.

Modelling software

The synthetic data are computed by a time-domain finite-difference code, made kindly available by Shell [5,6]. The code was written in 1993. The discretisation is sixth-order accurate in space and second-order in time, similar to the method used by Igel et al. [3]. It uses third-order Higdon absorbing boundary conditions [2]. Message passing is carried out via MPI. A master program takes care of the I/O and program control. This does not map in an optimal way to the Blue Gene hardware, because the nearest neighbour network topology cannot be exploited in this way. Still, the overall performance is much better than observed on standard Linux clusters.

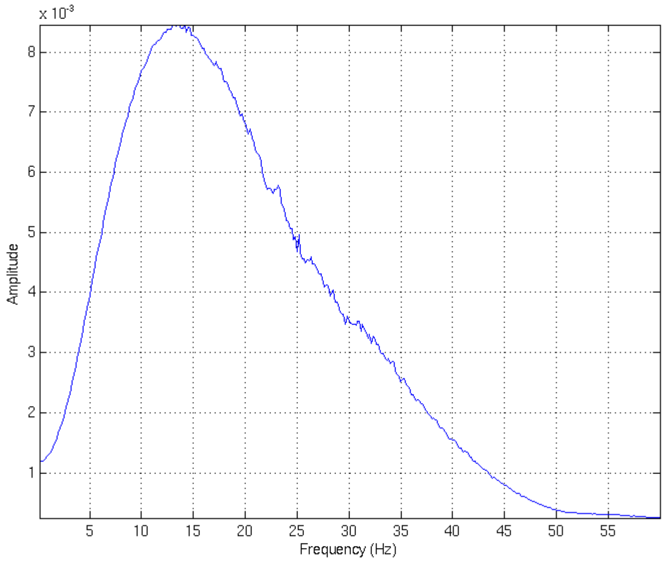

Figure 3: Average spectrum of the vertical displacement data.

Acquisition geometry

The final plan would be to complete a full line of 121 shots, with *x* between 2,012.5 m and 8,012.5 m and *y* at 11,375 m, using a 50 m interval. Receivers are placed on a grid with *x* between 1,000 and 9,000 m, and *y* between 11,375 m and 15,375 m at a 25 m spacing. The computations are symmetric with respect to line y = 11,375 m on which the sources are placed. The source can be explosive, at 8 m depth, or can be a horizontal or vertical force at the surface. In total, this will lead to four lines. Receivers in a deviated well are included, leading to a total of 51,846l. Six seconds of data at 2 ms are recorded for the displacements and their divergence. This produces SEGY files of 635 Mbytes for each shot and component. We may reduce the file size by sub sampling to a larger sampling interval.

The computation of one shot on a grid with 8 m spacing took about 3 days on 512 nodes. The peak frequency of the data was about 12 Hz. The 8-m spacing is too coarse by a factor two for realistic data frequencies: the peak frequency is about 6 Hz with reasonable accuracy up to 12 Hz. A single run on a 4-m grid with 5001×1251×1095 points took about 3 weeks on 1024 nodes, occupying one rack of the machine. This is too long for practical purposes. We will therefore concentrate on the 8 m runs, initially with a shot spacing of 250 m leading to 31 shots per line.

One shot panel computed on a 4 m grid is shown in Fig. 2. The data have been weighted linearly in time and are strongly clipped to bring out weak reflections and numerical dispersion. The average spectrum is shown in Fig. 3.

An indication of the accuracy can be obtained by comparing the result on the 4-m grid to that obtained for a grid spacing of 8 m. Because the wavelet for the latter was scaled accordingly, having a peak frequency at half the value used for the 4-m experiment, the data for the 4-m grid were filtered to obtain the same wavelet as used for the 8-m simulation. A comparison of traces for the horizontal and for the vertical displacement recorded on the surface at $x = 6$ km, is shown in Figure 4. Results compare fairly well. The strong surface wave is an exception and has a large error, caused by the fact that the explosive source is close to the free surface and by the approximation used at the free surface.

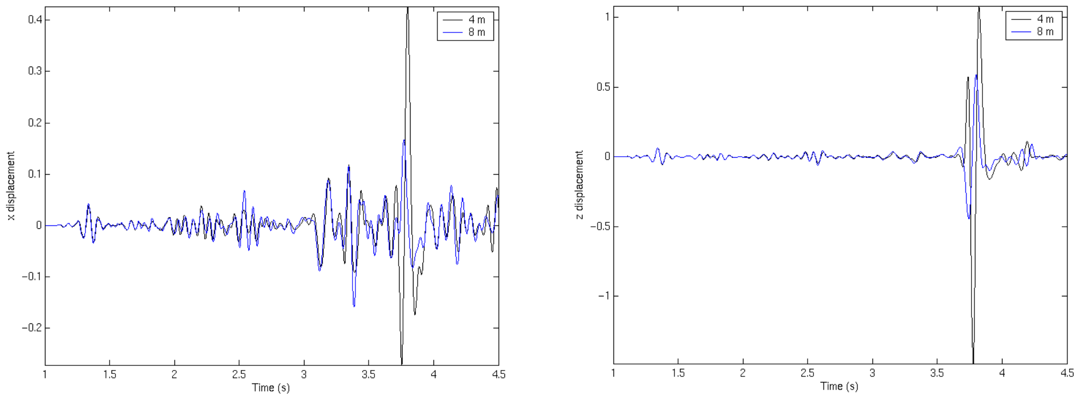

Figure 4. Comparison of traces for the horizontal (left) and vertical (right) displacement. The result for the 8-m grid spacing is shown in blue, the filtered result for the 4-m spacing in black.

Conclusions

We have constructed a two-dimensional anisotropic elastic model for generating three-dimensional synthetic data. The available hardware will only allow for a few shots on a grid that is sufficiently fine for realistic frequencies.

Acknowledgements

This project was initiated by the EAGE Research Committee. The authors thank the management of Shell for permission to use the modelling code, and Timothée Perdrizet of IFP for creating the index model. The computations were carried out on the LOFAR/University of Groningen Blue Gene/L system "Stella".

References

[1] Aminzadeh, F., Brac, J., and Kunz, T., 1997. 3-D Salt and Overthrust Models, Soc. of Expl. Geophysicists.

[2] Higdon, R. L., 1990. Radiation boundary conditions for elastic wave propagation. SIAM Journal on Numerical Analysis 27, 831–870.

[3] Igel, H., Mora, P. and Riollet, B., 1995. Anisotropic wave propagation through finite-difference grids. Geophysics 60, 1203–1216.

[4] LOFAR: http://www.lofar.org/

[5] Mulder, W.A., 1997. Experiments with Higdon's absorbing boundary conditions for a number of wave equations. Computational Geosciences 1, 85–108.

[6] Mulder, W.A., and Plessix, R.-E., 2002. Time- versus frequency-domain modelling of seismic wave propagation. Extended Abstract E015, 64th EAGE Conference & Exhibition, Florence, Italy.