3D Acoustic and Elastic Wave Modeling on a High Performance Computing System*

Suhas Phadke[1] and M. Ramanjulu[1]

Search and Discovery Article #40899 (2012)**
Posted March 31, 2012

*Adapted from extended abstract presented at GEO-India, Greater Noida, New Delhi, India, January 12-14, 2011

Research Laboratories Ltd., Pune, India

Abstract

Accurate modeling of seismic wave propagation in complex media is the key to many high-end seismic processing algorithms. The performance is achieved by parallel implementation of algorithms on large compute clusters. This article describes the implementation of 3D acoustic and elastic wave propagation algorithms on a high performance computing system. A MPI (Massage Passing Interface) implementation of complex modeling algorithms on EKA, Computational Research Laboratories' (CRL) high performance computing system is presented. Scalability and efficiency is achieved by proper structuring of the codes. Performance analysis is carried out for different sizes of the model. Application of seismic modeling is shown by generating seismic seismograms for some complex geological models.

Introduction

Seismic modeling is an area of significant research in the industry. Wave equation modeling is useful for understanding complex imaging problems, for processing and algorithm testing, and for AVO analysis. We can divide the main usage of seismic forward modeling into two parts, (1) predict seismic response (pre-acquisition modeling), and (2) understand seismic response (post-acquisition modeling). Pre-acquisition modeling can be used for survey design by understanding zones of illumination. Post-acquisition modeling can be useful for waveform inversion, AVO analysis, and for understanding the salt body locations, illumination of shadows, steep dips, gas clouds, artifacts, etc. The purpose of forward modeling is also to validate the geological model by comparing the synthetic data with the field data.

In the last two decades the power of a single CPU has increased by several orders of magnitude. But applications like seismic modeling need computational resources far greater than a sequential computer can provide. Parallel processing has proven to be a viable solution to improve performance, which uses the power of multiple computers connected together by a high-speed network.

Forward modeling, where the synthetic data is generated for a given Earth's model, is also a key step in the process of seismic inversion, where one tries to estimate the physical properties of the Earth. 80% to 90% of the computing time in an inversion algorithm is spent on generating synthetic data. Efficient parallel algorithms are therefore essential. Parallelism is the key to improve performance on any high performance computing system manufactured today.

The basic problem in theoretical seismology is to determine the wave response of a given model to the excitation of an impulse source by solving the wave equations. In scalar approximation, the acoustic wave equation is solved to evaluate the waveform by making medium velocities equal to compressional wave velocities. A more complete approach is to study the vector displacement field using full elastic wave equation for modeling both compressional and shear waves. In this formulation the mode conversions are automatically accounted for.

The mathematical model for elastic wave propagation in 3D heterogeneous media consists of coupled second order partial differential equations governing motions in x, y and z directions. Instead of solving second order coupled partial differential equations, we formulate them as a first order hyperbolic system (Virieux, 1986; Vafidis, 1988; Dai et al., 1996):

$$\frac{\partial Q}{\partial t} = A\frac{\partial Q}{\partial x} + B\frac{\partial Q}{\partial y} + C\frac{\partial Q}{\partial z}$$

Where Q is a vector comprising particle velocities and stress components and A, B, and C are matrices containing physical parameters of the medium. This formulation does not contain any derivatives of physical parameters. Thus, we need not calculate the gradients of physical parameters that may cause singularity in the numerical solution due to sharp changes in the physical properties.

The solution to the above partial differential equations is obtained by using finite difference methods. For this purpose the model is discretized into a number of grid points and the wavefield is calculated at each grid point as a function of time.

Parallel Implementation

The most important part of parallel programming is to map out a particular problem on a multiprocessor environment. The problem must be broken down into a set of tasks that can be solved concurrently. The choice of an approach to the problem decomposition depends on the computational scheme. In this study we have used the MacCormack scheme for finite differencing which is fourth order accurate in space and second order accurate in time. For this scheme, the calculation of the wavefield at a grid point of the advanced time level involves the knowledge of the wavefield at thirteen grid points of the current time level. Therefore, it is a thirteen point differencing star. Therefore, if we use a domain decomposition scheme for solving this problem second order neighbors will be involved in communication. Domain decomposition involves assigning subdomains of the computational domain to different processors and solving the equations for each subdomain concurrently. A checkerboard partitioning is applied here.

In checkerboard partitioning, domain is divided in all three directions creating smaller subdomains. In uniform checkerboard partitioning, all subdomains are of the same size. These subdomains have to be distributed among processors and no processor gets the complete plane. In order to calculate the wavefield at the grid points of the subdomain, at each time step, the required boundary grid points should be interchanged between the processors. In the case of checkerboard partitioning, the ghost point communication is in all three directions, i.e. each processor should exchange the boundary grid points with its six neighbors.

We have used MPI for the parallel implementation on a cluster of nodes. Each node comprises several cores. To set up the domain decomposition, we have used Cartesian topology approach of MPI. This helps in identifying the neighbors for the inter-processor communications during exchange of the wavefield along the subdomain boundaries. MPI Send-Receive function is used for this purpose. For all other parameter distributions usual MPI send, receive and broadcast functions have been used.

Performance Analysis

We performed the benchmark tests of the parallel 3D acoustic/elastic wave modeling algorithms for several problem sizes on EKA, CRL's high performance computing system (HPC). EKA is a HPC system with 1800 compute nodes, where each node is a dual quadcore processor. The nodes are connected by 20 GBps infiniband interconnect with less than 2μs latency. The peak and sustained performance of the system is 172 TFlops and 132.8 TFlops respectively.

The accuracy and efficiency of the method was tested by applying it to several models. Here we show the results of both 3D acoustic and elastic wave propagation for Marmousi-II model (Martin, Wiley and Marfurt, 2006). Originally this model is a 2D model, but we extended it in the third dimension by replicating the 2D model. The purpose was to test and benchmark the algorithms. Figure 1, shows the x-z slice of this complex model, where x denotes the horizontal axis and z denotes the vertical axis. Different colours represent different p-wave velocities in the model. Shear wave velocity and density is also assigned to each grid point of the model.

In this paper we show the results of 3D acoustic and elastic wave modeling for a problem size of 1400 X 400 X 800. Many more results and examples will be shown during presentation. For the model shown in Figure 1, Δx and Δz are 6.25 m and 5.0 m respectively and the time step is 0.5 msec. The source is at 1.25 km and the data is recorded by 580 receivers located from 1.35 km to 8.6 km. The nearest offset is 100 m and the farthest offset is 7250 m. 8000 time steps are required in order to produce synthetic data for 4.0 sec. Figure 2 shows the synthetic seismograms obtained by 3D acoustic and elastic wave modeling algorithm for Marmousi-II model. This model is a good representation of a marine model as the topmost layer is a 500 m thick water layer. Source and receivers are located in the water layer. As shear waves do not propagate through the water layer, only pressure waves are recorded by the receivers placed in the water layer. Therefore both the acoustic and elastic wave seismograms in Figure 2 appear similar. Since the model is quite complex we do not explicitly see any intrabed and converted wave multiples, although they are also present in the seismograms.

Table 1 and Table 2 show the performance of the acoustic and elastic wave propagation algorithms. A minimum of 32 cores were needed to run this size of the model. Therefore all the other runs are compared with the 32 core run. The compute times are compared for 1000 time steps only. As the number of cores increased from 32 to 512, a sixteen fold increase in the number of cores, the compute time for acoustic wave modeling reduced by a factor of 10, approximately 64% efficiency. Similarly for elastic wave modeling, the reduction in compute time is by a factor of 11 and the efficiency is 68%. From the tables one can easily observe that a large amount of compute time is necessary to produce one shot gather.

Discussion and Conclusions

In this paper the implementation of 3D acoustic and elastic wave modeling algorithms is carried out on a high performance computing system. The codes for these parallel implementations have been written using MPI message passing libraries. We have discussed an example for showing performance of the algorithm. Performance can be improved by proper structuring of the code. Reverse Time Migration (RTM), which is a state-of-the-art migration technique for imaging, makes use of parallel implementation of 3D acoustic wave modeling algorithm. There is a need to improve performance and reduce compute time for making RTM a routine seismic processing algorithm. Forward modeling is also an integral part of any inversion algorithm. Waveform inversion technique, which determines an optimal subsurface model by comparing observed seismograms with the synthetic seismogram, has to compute synthetic seismograms thousands of times. The presentation will focus on both, the modeling algorithms and their application in seismic data processing and inversion.

Acknowledgment

Authors wish to thank Computational Research Laboratories Ltd., Pune for providing computational facility on EKA and permission to publish this work.

References

Vafidis, A., 1988, Supercomputer finite difference methods for seismic wave propagation: Ph.D. Thesis, University of Alberta, Edmonton, Canada.

Virieux, J., 1986, P-SV wave propagation in heterogeneous media: velocity stress finite difference method: Geophysics, v. 51, p. 889-890.

Dai, N., A. Vafidis, and E.R. Kanasewich, 1996, Seismic migration and absorbing boundaries with a one way wave system for heterogeneous media: Geophysical Prospecting, v. 44, p. 719-739.

Martin, G.S., R. Wiley, and K.J. Marfurt, 2006, Marmousi2: An elastic upgrade for Marmousi: The Leading Edge, p. 156-166.

Mitchell, A.R., and D.F. Griffiths, 1981, The finite difference method in partial differential equations: John Wiley & Sons Inc., 284 p.

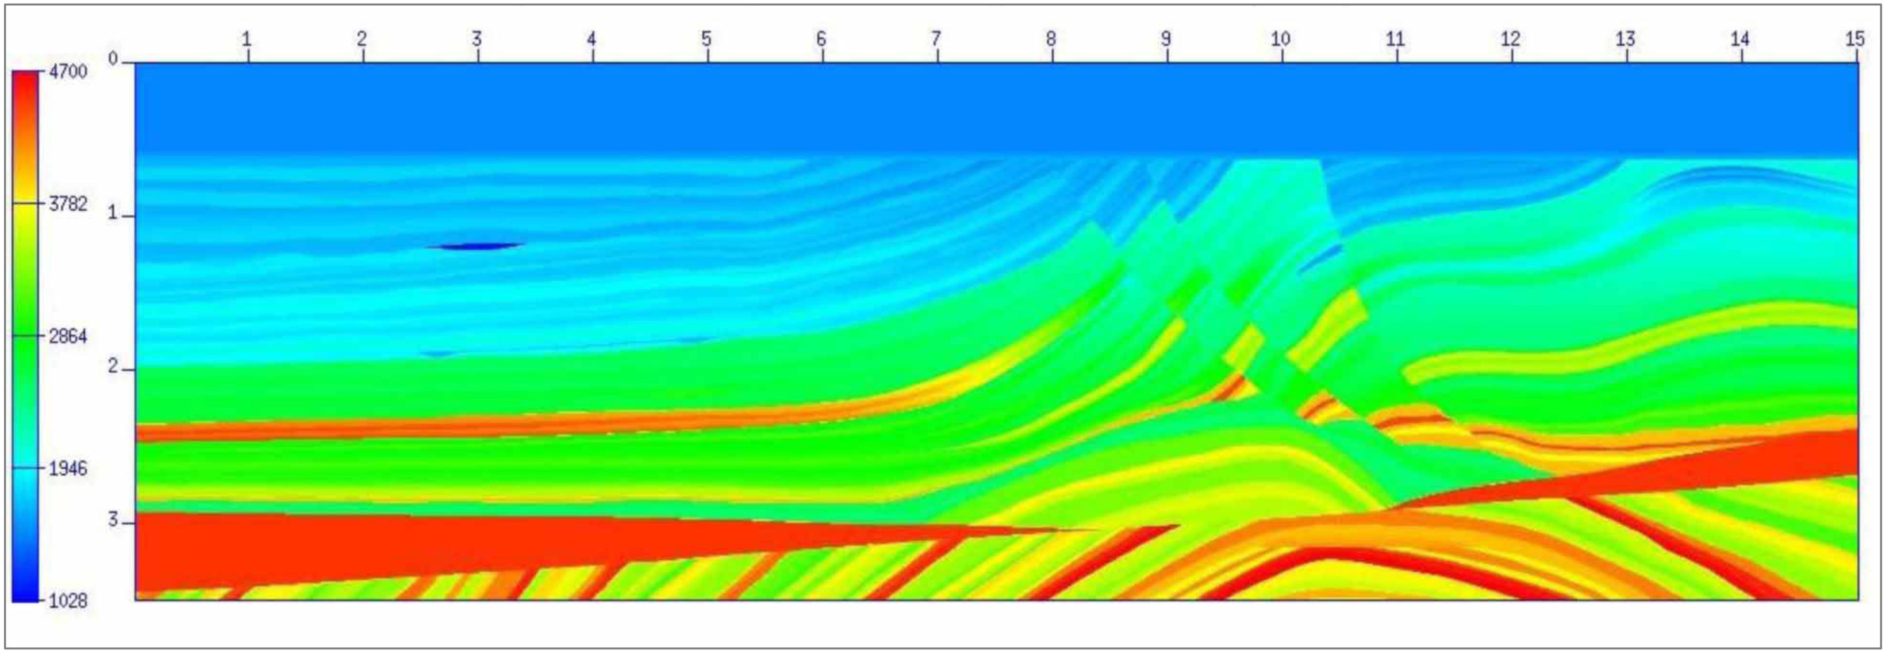

Figure 1. A schematic of Marmousi-II velocity model. The colour bar on left represents p-wave seismic velocities in m/sec.

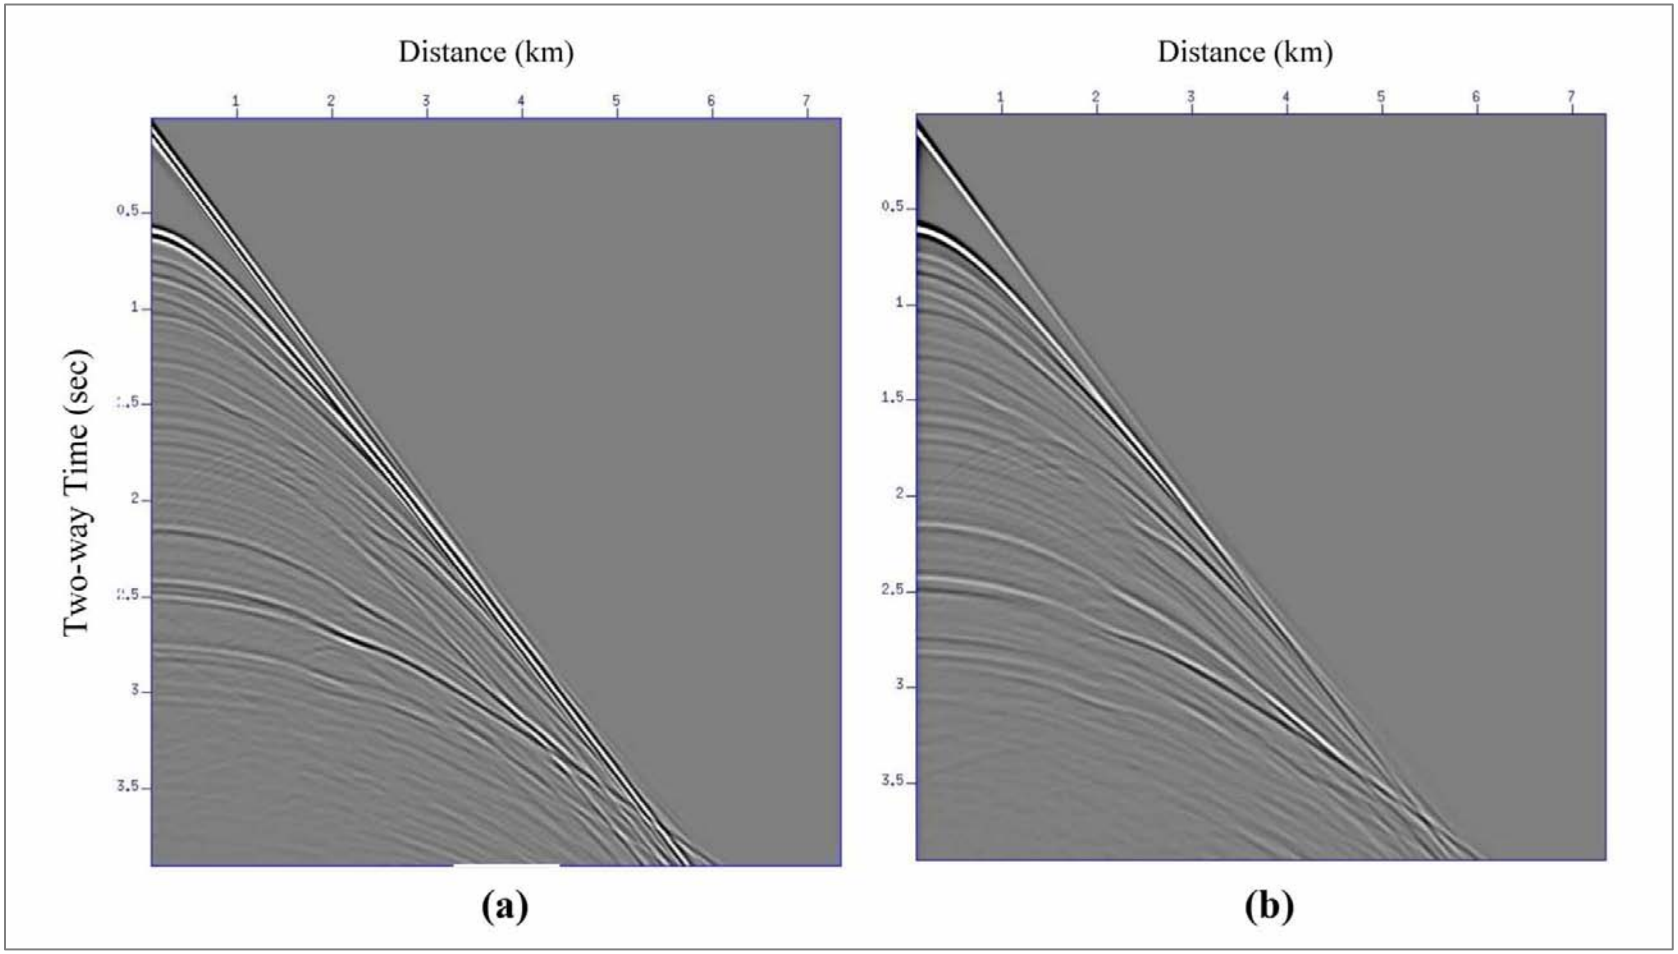

Figure 2. Synthetic seismograms obtained from 3D acoustic and elastic wave modeling algorithm. (a) Acoustic wave synthetic seismograms, and (b) Elastic wave synthetic seismograms.

Grid Size	#Cores	Time for 1000 time steps(sec)	Relative speed up over 32 cores	Relative efficiency over 32 cores
1400×400×800	32	6240	1	100.00%
1400×400×800	64	3323	1.88	94.00%
1400×400×800	128	1799	3.47	86.75%
1400×400×800	256	1036	6.02	75.25%
1400×400×800	512	606	10.3	64.37%

Table 1. Performance of Acoustic wave modeling algorithm as a function of number of cores.

Elastic Forward Modelling Results				
Grid Size	#Cores	Time for 1000 time steps(sec)	Relative speed up over 32 cores	Relative efficiency over 32 cores
1400×400×800	32	21013	1	100.00%
1400×400×800	64	10672	1.97	98.44%
1400×400×800	128	5538	3.79	94.85%
1400×400×800	256	3131	6.71	83.89%
1400×400×800	512	1914	10.98	68.61%

Table 2. Performance of the Elastic wave modeling algorithm as a function of number of cores.