3D Finite-Difference Seismic Migration with Parallel Computers *

Curtis C. Ober † Robert Gjertsen † Susan Minkoff † David E. Womble †

Abstract

The ability to image complex geologies such as salt domes in the Gulf of Mexico and thrusts in mountainous regions is essential for reducing the risk associated with oil exploration. Imaging these structures, however, is computationally expensive as datasets can be terabytes in size. Traditional ray-tracing migration methods cannot handle complex velocity variations commonly found near such salt structures. Instead we use the full 3D acoustic wave equation, discretized via a finite difference algorithm. We reduce the cost of solving the paraxial wave equation by a number of numerical techniques including the method of fractional steps and pipelining the tridiagonal solves. The imaging code, Salvo, uses both frequency parallelism (generally 90% efficient) and spatial parallelism (65% efficient). Salvo has been tested on synthetic and real data and produces clear images of the subsurface even beneath complicated salt structures.

1 Introduction

To obtain information about the earth's interior, oil and gas companies perform many thousands of acoustic experiments — setting off "shots" and recording the portions of the propagated waves which reflect off internal heterogeneities and return to the surface. One common seismic processing technique takes this recorded wave data and converts it from the time to the depth domain to produce an image of the subsurface. This procedure, migration, can take weeks or even months of computer time if a large number of shots (millions) must be processed.

Oil companies traditionally use ray-tracing (Kirchhoff) methods for seismic processing of new prospects. These schemes are economical but cannot account for multiple travel paths or strong velocity variations in regions of complex geology. Salt intrusions are an example of a complex geologic formation commonly found above oil reservoirs in the Gulf of Mexico. Clearly, some of the most interesting prospects for hydrocarbon recovery are regions with complex geology. To overcome the limitations of ray tracing, we have implemented a finite-difference algorithm for solving the acoustic wave equation in 3D. Because finite-difference methods are more expensive than ray-tracing, we have implemented the migration in parallel. The wave equation is transformed via the paraxial approximation and operator splitting. A subsequent corrective filter compensates for errors introduced in the approximation. The algorithm uses both frequency and spatial parallelism to produce good 3D images of the subsurface in reasonable time. In the remainder of this paper we cover the basic equations governing the migration, the numerical approximations we make, the parallel solution strategy, and some experimental results.

*This work was supported by the United States Department of Energy under Contract DE-AC04-94AL85000, and by DOE's Office of Mathematical, Information and Computational Sciences.

†Sandia National Laboratories, Albuquerque, NM.

DISCLAIMER

DISCLAIMER

Portions of this document may be illegible electronic image products. Images are produced from the best available original document.

2 Governing Equation

To model wave propagation within the earth's interior, we begin with the scalar wave equation,

$$\frac{\partial^2 P}{\partial x^2} + \frac{\partial^2 P}{\partial y^2} + \frac{\partial^2 P}{\partial z^2} = \frac{1}{v^2}\frac{\partial^2 P}{\partial t^2}, \tag{1}$$

Here $v(x,y,z)$ is the acoustic velocity of the medium, and $P(x,y,z,t)$ is the pressure wavefield. Data collected at the surface, $P(x,y,0,t)$, will be downward propagated to illuminate reflector geometry within the subsurface. The scalar wave equation in 3D, Eq. (1), is time consuming to solve. The solution time can be reduced by transforming the hyperbolic equation (time domain) to a parabolic equation in depth through the paraxial approximation [2](p.54).

$$\frac{\partial P}{\partial z} = \left\{ \pm \frac{i\omega}{v} \left[1 + \frac{v^2}{\omega^2} \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right) \right]^{1/2} \right\} P \tag{2}$$

is the paraxial wave equation, and P is now a function of (x,y,z,ω). The positive and negative signs correspond to upcoming and downgoing wave fields respectively. Surface pressure is used as an initial condition for solving the resulting parabolic equation down to the next depth level. Solving Eq. (2) for every frequency to be included in the solution, we will likely have hundreds of paraxial equations to solve.

The square-root operator of the paraxial wave equation, Eq. (2), is difficult to estimate numerically. Therefore, this operator is approximated by an optimized series which has its origin in a continued fraction expansion [2](p.84) [8](p.513). The continued fraction expansion can be represented by ratios of polynomials [6], and the polynomial coefficients can be optimized for propagation angle [4]. With these approximations, we write the paraxial wave equation as

$$\frac{\partial P}{\partial z} = \pm \frac{i\omega}{v} \left[1 + \sum_{\ell=1}^{m} \frac{\alpha_\ell S}{1 + \beta_\ell S} \right] P, \qquad \text{where} \qquad S = S_x + S_y = \frac{v^2}{\omega^2}\frac{\partial^2}{\partial x^2} + \frac{v^2}{\omega^2}\frac{\partial^2}{\partial y^2}, \tag{3}$$

and α_ℓ and β_ℓ are the expansion coefficients derived by Lee and Suh [4]. Commonly one term of the series is retained ($m = 1$), and propagation angles of up to 65 degrees can be accurately determined.

3 Method of Solution

To solve Eq. (3) by a finite-difference technique requires solving a banded matrix. The size of this matrix is on the order of $N^2 \times N^2$ where N is the number of grid points in either the x or y directions. Clearly, a direct solution of such a large matrix is prohibitive. To reduce the computational cost, each term in Eq. (3) is solved separately by the method of fractional steps [3]. For instance, the first equation

$$\frac{\partial P_0}{\partial z} = \pm \frac{i\omega}{v(x,y,z)} P_0, \tag{4}$$

has the exact solution $P_0(x,y,z+\delta z,\omega) = P_0(x,y,z,\omega)\exp[\pm(i\omega/v)\Delta z]$. The second equation is given by

$$\frac{\partial P_1}{\partial z} = \pm \frac{i\omega}{v} \frac{\alpha_1 S}{1 + \beta_1 S} P_1. \tag{5}$$

Solving Eq. (5) is computationally intensive. To reduce this expense, another operator separation is applied in the x and y directions [1] so that efficient tridiagonal solves can be performed via the Thomas Algorithm [3]. To convert the operator, S, to a linear combination of S_x and S_y, we write

$$\frac{\partial P_1}{\partial z} = \pm\frac{i\omega}{v}\left[\frac{\alpha_1 S_x}{1+\beta_1 S_x} + \frac{\alpha_1 S_y}{1+\beta_1 S_y} + \frac{-2\alpha_1\beta_1 S_x S_y}{(1+\beta_1 S)(1+\beta_1 S_x)(1+\beta_1 S_y)}\right.$$

$$\left. + \frac{-\alpha_1\beta_1^2 (S_x S_x S_y + S_x S_y S_y)}{(1+\beta_1 S)(1+\beta_1 S_x)(1+\beta_1 S_y)}\right] P_1 \tag{6}$$

or, dropping cross-terms,

$$\frac{\partial P_1}{\partial z} \approx \pm\frac{i\omega}{v}\left[\frac{\alpha_1 S_x}{1+\beta_1 S_x} + \frac{\alpha_1 S_y}{1+\beta_1 S_y}\right] P_1. \tag{7}$$

The operators in Eq. (7) are once again split by the method of fractional steps, and solved in sequence

$$\frac{\partial P_1}{\partial z} = \pm\frac{i\omega}{v}\frac{\alpha_1 S_x}{1+\beta_1 S_x} P_1 \quad \text{and} \quad \frac{\partial P_1}{\partial z} = \pm\frac{i\omega}{v}\frac{\alpha_1 S_y}{1+\beta_1 S_y} P_1. \tag{8}$$

Operator splitting is very similar to approximate factorization (AF) or ADI schemes where the x- and y-directions are separated to produce tridiagonal systems of equations. The method of fractional steps (MOFS) can be shown to be equivalent to the ADI scheme if the operators in Eqs. (8) commute and we can neglect a term of order Δz^2. However, MOFS only requires one evaluation of the right-hand-side terms where ADI schemes require two evaluations of the right-hand-side. Also, the MOFS scheme for the paraxial wave equation only requires one level of storage ($n_x \times n_y$) for the pressure wavefield rather than two as for AF and ADI schemes. Thus to maintain low memory requirements, MOFS is the method of choice. After operator splitting and approximations, the paraxial wave equation reduces to

$$\frac{\partial P}{\partial z} = \pm\frac{i\omega}{v}\left[1 + \frac{\alpha_1 S_x}{1+\beta_1 S_x} + \frac{\alpha_1 S_y}{1+\beta_1 S_y}\right] P. \tag{9}$$

Neglecting the two cross-terms in Eq. (6) introduces a phase error which is maximum along the lines $y = \pm x$. Li [5] compares Eqs. (2) and (9) and attempts to compensate for <u>all</u> the errors between them by using the phase error,

$$E = \sqrt{1 + S_x + S_y} - \left[1 + \sum_{\ell=1}^{m} \frac{\alpha_\ell S_x}{1+\beta_\ell S_x} + \sum_{\ell=1}^{m} \frac{\alpha_\ell S_y}{1+\beta_\ell S_y}\right]. \tag{10}$$

Applying the error operator to the wavefield, P,

$$\frac{\partial P}{\partial z} = \left[\frac{i\omega}{v_a(z)} E\right] P, \tag{11}$$

we can partially correct for the errors induced by approximation. In the expression above, $v_a(z)$ is the average acoustic velocity in any depth plane. The solution to this equation is $P(z+\Delta z) = P(z)\exp[E\Delta z(iw/v_a)]$ and is applied as an additional step in the method of fractional steps. We note that if the acoustic velocity is only a function of depth (i.e., a layered medium), the Li filter exactly corrects for all the phase errors, including the approximation errors resulting from the difference between Eqs. (2) and (9).

4 Parallelization

There are two types of parallelism in the seismic imaging algorithm described above: frequency parallelism and spatial parallelism. In frequency parallelism, each processor migrates a subset of the frequencies, and there is relatively little communication required for this. The velocity model is distributed to all processors at the beginning of a migration step, and the image is constructed by summing pressure fields across the processors. The limitations of frequency parallelism are that the number of processors must be smaller than the number of frequencies being processed and that each processor must store the full velocity model. These problems are mitigated by spatial parallelism in which the (x, y) domain is also distributed among processors. A new problem arises — the tridiagonal solves in the x and y directions must be parallelized.

It is difficult to parallelize the solution of a single tridiagonal system, but this difficulty is offset because there are many such systems. In our algorithm, Salvo, we set up a pipeline. In the first stage of the pipeline, processor one starts on several tridiagonal solves. In the second stage of the pipeline, processor two continues the first set of tridiagonal solves, while processor one starts a second set of tridiagonal solves. This process continues until all processors are busy. The optimal number of tridiagonal solves in a set is determined by minimizing the overhead time associated with communication and idle time (see [7]).

5 Results and Conclusions

We tested Salvo with several synthetic datasets, as well as on real data. As an example of a Salvo migration, we show the 3D SEG/EAGE salt model. Figure 1(a) shows a corner–cut view of the velocity model. The grayscale colormap indicates the speed of sound in a region; the lighter a region, the higher the speed. The white region is the salt dome. Figure 1(b) shows the same corner–cut view of the image produced by Salvo. The image is $600 \times 600 \times 210$. It is a stack of 45 shots, each processed on a $200 \times 200 \times 210$ grid with a surface grid fully populated with receivers. A total of 511 frequencies were migrated.

To test the computational performance of Salvo, the sample impulse problem was run on the Intel Paragon. Spatial parallelism was tested by setting $p_\omega = 1$ and varying the number of processors in the x and y directions, $p_x \times p_y$ and the number of grid points, $n_x \times n_y$ (scaled-problem size). For frequency parallelism, we varied the number of processors in the ω direction, p_ω, while keeping $p_x \times p_y = 1$ and the number of frequencies, n_ω, constant.

Timings for the sample impulse run are shown in Table 1. The spatial parallelism is efficient when the pipeline is fully utilized (3×3 nodes), remaining nearly constant at 66%. Frequency parallelism is very efficient, remaining close to 90% for most problems. The drop off in efficiency for the 64-node run is due to too little work for each processor.

References

[1] D. L. Brown, *Applications of Operator Separation in Reflection Seismology*, Geophysics, 48 (1983), pp. 288–294.

[2] J. F. Claerbout, *Fundamentals of Geophysical Data Processing with Applications to Petroleum Propecting*, Blackwell Scientific Publications, 1985.

[3] C. Fletcher, *Computational Techniques for Fluid Dynamics Vol. I*, Springer-Verlag, Berlin, 1988.

[4] M. W. Lee and S. Y. Suh, *Optimization of One-Way Wave Equations*, Geophysics, 50 (1985), pp. 1634–1637.

[5] Z. Li, *Compensating Finite-Difference Errors in 3-D Migration and Modeling*, Geophysics, 56 (1991), pp. 1650–1660.

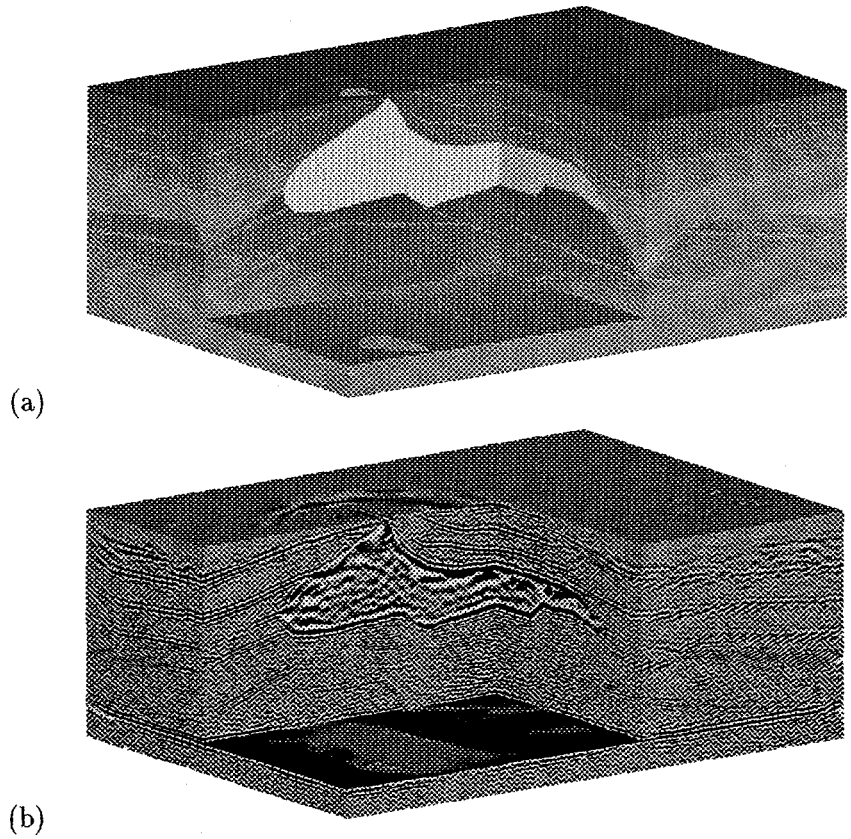

FIG. 1. *A corner–cut view of the SEG/EAGE 3D salt model and the same corner cut view of the Salvo image produced by processing 45 shots.*

Spatial Parallelism			Frequency Parallelism		
$p_x \times p_y$	$n_x \times n_y \times n_\omega$	Efficiency (%)	p_ω	$n_x \times n_y \times n_\omega$	Efficiency (%)
1×1	$101 \times 101 \times 32$	100.0	1	$101 \times 101 \times 128$	100.0
3×3	$301 \times 301 \times 32$	71.6	2	$101 \times 101 \times 64$	100.0
4×4	$401 \times 401 \times 32$	71.3	4	$101 \times 101 \times 32$	101.6
5×5	$501 \times 501 \times 32$	69.5	8	$101 \times 101 \times 16$	98.6
6×6	$601 \times 601 \times 32$	68.7	16	$101 \times 101 \times 8$	98.6
7×7	$701 \times 701 \times 32$	67.2	32	$101 \times 101 \times 4$	92.1
8×8	$801 \times 801 \times 32$	66.7	64	$101 \times 101 \times 2$	85.9

TABLE 1

Timings for a sample impulse problem for spatial and frequency on the Intel Paragon.

[6] Z. Ma, *Finite-Difference Migration with Higher Order Approximation*, in 1981 Joint Mtg. Chinese Geophys. Soc. and Society of Exploration Geophysicists, Society of Exploration Geophysicists, 1981.

[7] C. Ober, R. Oldfield, D. Womble, and V. J.P., *Seismic Imaging on the Intel Paragon*, J. Computers & Math. with Applic., (1998). To Appear.

[8] O. Yilmaz, *Seismic Data Processing, Investigations in Geophysics No. 2*, Society of Exploration Geophysicists, P.O. Box 702740, Tulsa, OK 74170-2740, 1987.