ITERATIVE METHODS FOR 3D IMPLICIT FINITE-DIFFERENCE MIGRATION USING THE COMPLEX PADÉ APPROXIMATION

C. A. N. Costa, I. S. Campos, J. C. Costa, F. A. Silva Neto, J. Schleicher, and A. Novais

email: *js@ime.unicamp.br*

keywords: *3D Seismic Imaging, wave-equation, sparse-solvers*

ABSTRACT

Conventional implementations of 3D finite-difference (FD) migration use splitting techniques to accelerate performance and save computational cost. However, such techniques are plagued with numerical anisotropy that jeopardizes correct positioning of dipping reflectors in the directions not used for the split operators. We implemented a 3D downward continuation FD migration without splitting in the space coordinates using a complex Padé approximation and implicit finite differences. In this way, the numerical anisotropy is eliminated at the expense of a computationally more intensive solution of a large banded linear system. We compare the performance of the iterative stabilized biconjugate gradient (BICGSTAB) and the multifrontal massively parallel direct solver (MUMPS). It turns out that the use of the complex Padé approximation not only stabilizes the solution, but also acts as an effective preconditioner for the BICGSTAB algorithm, reducing the number of iterations as compared to implementation using the real Padé expansion. As a consequence, the iterative BICGSTAB method was more efficient than the direct MUMPS method when solving for a single term in the Padé expansion. The results of both algorithms, as evaluated computing the migration impulse response in the SEG/EAGE salt model, are of comparable quality.

INTRODUCTION

Efficient implementations of 3D FD migration use the splitting technique along horizontal coordinates to avoid the solution of large banded linear systems (Claerbout, 1985) at each downward continuation step. Splitting the migration operator along inline and crossline directions reduces 3D migration to a sequence of 2D downward continuation steps making the algorithm highly efficient. As a drawback splitting introduces numerical anisotropy which can damage the image of reflectors (Brown, 1983). Several strategies are used to remedy this problem usually assuming a homogeneous medium and using wavefield interpolation to handle lateral velocity variations (Li, 1991). Another idea is to increase the number of splitting directions (Ristow and Rühl, 1997). However, such correction methods add again to the cost and introduce additional errors and artifacts. Exploring another idea to avoid the expensive direct solution of these large systems, Cole (1989) investigated the use of iterative methods to implement FD migration without splitting and evaluated the performance of overrelaxation, Jacobi, and Gauss-Seidel methods (Iserles, 1996). He reported the poor conditioning of the linear system matrix for low frequencies which degraded the performance of iterative solvers to a prohibitively high computational cost. Along a similar line, Nichols (1991) investigated the performance of the conjugate gradient method (CG) in dependence on the choice of the initial value for preconditioning. He found a degradation of performance of CG for low frequencies but reported that the performance improves when the propagation of energy associated with high angles and evanescent modes is attenuated.

In this respect, it is interesting to note that Amazonas et al. (2007) improved the quality of migrated images by representing the downward continuation operator using the complex Padé approximation (Millinazzo et al., 1997), which is essentially a means of attenuating the propagation of evanescent and high-angle wave modes. Moreover, new iterative methods to solve complex linear system are available that allow to compute the solution of very large linear systems. For example, the stabilized biconjugate gradient (BICGSTAB) method can solve poor conditioned complex linear systems more efficiently than CG (van der Vorst, 1992). Another development in this area is the multifrontal massively parallel direct solver (MUMPS) which can efficiently solve large complex banded linear systems using Gaussian decomposition (Amestoy et al., 2001, 2006). In this paper, we investigate whether the limitations observed by Nichols (1991) can be overcome when using the complex Padé approximation in the implementation of 3D downward continuation by finite differences without splitting. Our numerical experiments in homogeneous media and in the SEG/EAGE salt model indicate that the complex Padé approximation is an effective preconditioner for 3D FD migration, so that the iterative BICGSTAB approach outperforms the massively parallel MUMPS solver when using a single term of the Padé expansion. If more terms are needed to better model high dip events the direct solver is required.

THEORY

Implicit finite-difference downward continuation

The one-way wave equation for downward continuation in the space-frequency domain is

$$\frac{\partial P(\mathbf{x},\omega)}{\partial z} = -\frac{i\omega}{c(\mathbf{x})}\sqrt{1+D}\,P(\mathbf{x},\omega)\,, \tag{1}$$

where $\mathbf{x} = (x, y, z)$ is the position with x and y denoting the coordinates along inline and crossline directions and z denoting depth, ω is the angular frequency. Moreover, $P(\mathbf{x},\omega)$ is the wavefield spectrum and $c(\mathbf{x})$ represents the wave speed. Finally, D represents the differential operator $D = \frac{c(\mathbf{x})^2}{\omega^2}(D_x^2 + D_y^2)$, with $D_x \equiv \frac{\partial}{\partial x}$ and $D_y \equiv \frac{\partial}{\partial y}$.

The square-root operator in equation (1) can expanded using the complex Padé series (Millinazzo et al., 1997; Amazonas et al., 2007):

$$\frac{\partial P}{\partial z} = -\frac{i\omega}{c(\mathbf{x})}\left[C_0 + \sum_{n=1}^{N}\frac{A_n D}{1+B_n D}\right]P(\mathbf{x},\omega)\,, \tag{2}$$

where N indicates the order of the Padé expansion and C_0, A_n, and B_n are the complex coefficients

$$\begin{aligned} C_0 &= e^{i\theta/2}\left[1+\sum_{n=1}^{N}\frac{a_n\left(e^{-i\theta}-1\right)}{1+b_n\left(e^{-i\theta}-1\right)}\right], \\ A_n &= \frac{a_n e^{-i\theta/2}}{\left[1+b_n\left(e^{-i\theta}-1\right)\right]^2}, \\ B_n &= \frac{b_n e^{-i\theta}}{1+b_n\left(e^{-i\theta}-1\right)}. \end{aligned} \tag{3}$$

Here, θ indicates the angle of rotation for the branch cut of the square root in the complex plane, and a_n and b_n are the real Padé expansion coefficients (Bamberger et al., 1988)

$$a_n = \frac{2}{2N+1}\sin^2\frac{n\pi}{2N+1} \qquad \text{and} \qquad b_n = \cos^2\frac{n\pi}{2N+1}\,. \tag{4}$$

The constant C_0 does not depend on D and is actually an approximation to one that gets better with an increasing number of terms N used in the expansion. Therefore, we can set $C_0 = 1$.

The numerical solution of equation (2) is obtained by solving a sequence of differential equations (Claerbout, 1985). At each continuation step Δz, the solution is computed recursively for each term of the

expansion, using the solution of the previous step as initial condition for the next one. The first equation to be solved is

$$\frac{\partial P(\mathbf{x},\omega)}{\partial z} = -i\frac{\omega}{c(\mathbf{x})}P(\mathbf{x},\omega)\,, \tag{5}$$

followed by the N equations

$$\frac{\partial P(\mathbf{x},\omega)}{\partial z} = -i\frac{\omega}{c(\mathbf{x})}\left[\frac{A_n D}{1+B_n D}\right]P(\mathbf{x},\omega) \qquad (n=1,\ldots,N)\,. \tag{6}$$

This algorithm converges for small Δz.

Using the Crank-Nicolson finite-differences scheme (Iserles, 1996) to approximate the derivative operators, equations (6) are reduced to the linear systems

$$\left(\mathbf{I}+\alpha^{+}\mathbf{D}\right)\mathbf{P}^{k+1} = \left(\mathbf{I}+\alpha^{-}\mathbf{D}\right)\mathbf{P}^{k}\,, \tag{7}$$

where

$$\alpha^{\pm} \equiv B_n \pm i\frac{\omega\Delta z}{c}A_n\,, \tag{8}$$

$$\mathbf{D} \equiv \left(\frac{c^2}{\omega^2}\right)\left[\frac{\mathbf{D}_x}{(\Delta x)^2}+\frac{\mathbf{D}_y}{(\Delta y)^2}\right]\,. \tag{9}$$

Here, $\mathbf{I}$ is the identity matrix, $\mathbf{D}_x$ and $\mathbf{D}_y$ are the finite-differences matrices representing the derivatives along the x and y coordinates, and $\mathbf{P}^k$ and $\mathbf{P}^{k+1}$ represent the wavefield at the grid at successive depth levels. The matrix $\mathbf{D}$ has dimension $N_x \times N_y$. Though sparse, this matrix can have a large band. This increases the computation cost of solving system (7) using direct methods. An efficient and frequently used method to reduce this cost is splitting along inline and crossline directions,

$$\left(\mathbf{I}+\alpha^{+}\left(\frac{c}{\omega\Delta y}\right)^2\mathbf{D}_x\right)\left(\mathbf{I}+\alpha^{+}\left(\frac{c}{\omega\Delta x}\right)^2\mathbf{D}_y\right)\mathbf{P}^{k+1} = \left(\mathbf{I}+\alpha^{-}\left(\frac{c}{\omega\Delta y}\right)^2\mathbf{D}_x\right)\left(\mathbf{I}+\alpha^{-}\left(\frac{c}{\omega\Delta x}\right)^2\mathbf{D}_y\right)\mathbf{P}^{k}\,. \tag{10}$$

The solution of this linear system is a sequence of tridiagonal linear systems each one of dimension N_x and N_y. Unfortunately, this approximation introduces numerical anisotropy through the incorrect mixed-derivatives term. Figure 1 compares the wavefield for a homogeneous medium, computed with splitting along the inline and crossline directions according to equation (10), to the one obtained without splitting, equation (7), at the same depth. The red circles in these images indicate the correct position. The noncircular format of the result of splitting is the consequence of numerical anisotropy.

To reduce numerical anisotropy, Li (1991) proposed a correction in the frequency-wavenumber domain which requires wavefield interpolation in laterally heterogeneous media and Ristow and Rühl (1997) used multiway splitting. However, such correction methods add to the computational cost and introduce additional errors and migration artifacts. For this reason, we are interested in alternative methods to solve system (7) that do not make use of splitting techniques.

Iterative methods

The sparseness of linear system (7) makes it suitable for iterative methods (Iserles, 1996). Cole (1989) investigated the performance of some iterative methods to solve equation (7) using the first order real Padé expansion. He noticed the rather slow convergence, or even divergence, of methods like overrelaxation, Jacobi, and Gauss-Seidel, when applied for frequencies below $\omega^2 = 8c^2b_n/(\Delta x)^2$. Nichols (1991) used the conjugate gradient method (CG) to solve equation (7) because of its better theoretical convergence properties, but also reported poor performance of this method for low frequencies. However, he observed some improvement in the convergence speed of CG when evanescent modes are filtered out. Unfortunately, the performance degrades rapidly for ill conditioned matrices, because CG solves the normal equations.

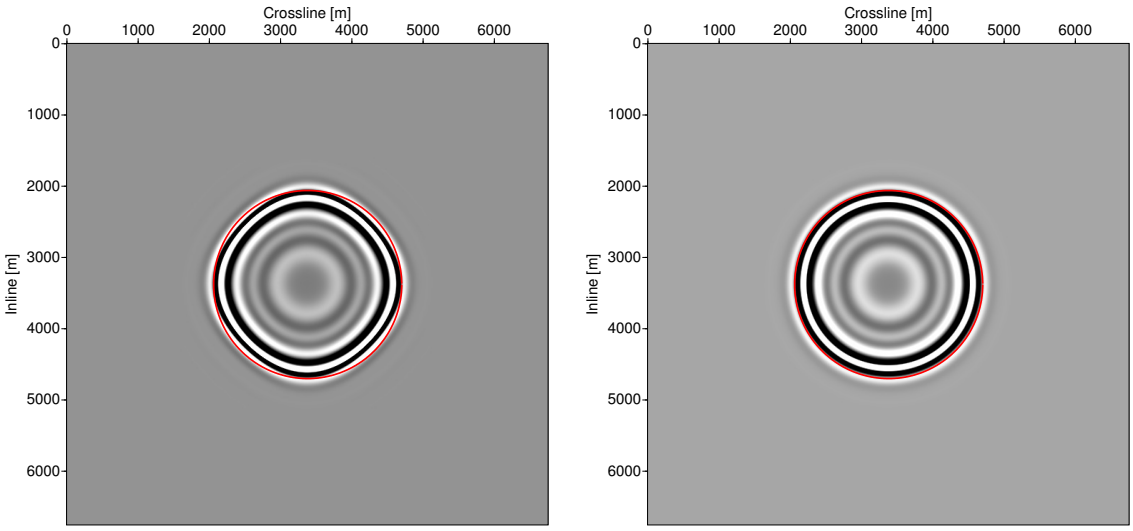

Figure 1: FD impulse response in a homogeneous medium. Left: splitting along the inline and crossline directions; Right: no splitting.

Since the days of Cole (1989) and Nichols (1991), iterative inversion techniques have improved. We investigate whether the limitations of conventional iterative methods as reported in the cited works still hold for more modern techniques. Taking into account the observation of Nichols (1991) that the performance of CG improves when the propagation of energy associated with high angles and evanescent modes is attenuated, we combine the complex Padé expansion (Amazonas et al., 2007) with the stabilized biconjugate gradient (BICGSTAB) algorithm (van der Vorst, 1992) to implement downward continuation without splitting. We chose BICGSTAB because of its better convergence properties than conventional CG for complex matrices.

Conditioning The performance of iterative methods improves when the matrix to be inverted is strictly diagonally dominant (Iserles, 1996). For a matrix $\mathbf{A}$ this means

$$|A_{jj}| \geq \sum_{k=1}^{N_A} |A_{jk}|_{j \neq k} \qquad (j = 1, ..., N_A) \,, \tag{11}$$

where N_A denotes the dimension of matrix $\mathbf{A}$. For a homogeneous medium we can verify when the matrix on the left side of equation (7) satisfies this condition. Its diagonal entries are

$$A_{jj} = \left(\frac{\omega \Delta x}{c} \right)^2 - 4 \left(B_n + i \frac{\omega \Delta x}{2c} A_n \right) \,, \tag{12}$$

and the only four nonzero off diagonal entries in each row are all equal to

$$A_{jk} = B_n + i \frac{\omega \Delta x}{2c} A_n \qquad (j \neq k) \,. \tag{13}$$

Thus, this matrix is diagonally dominant if

$$\omega \geq \omega_L = \frac{2c}{\Delta x} \left[- \operatorname{Im} \{A_n\} + \sqrt{(\operatorname{Im} \{A_n\})^2 + 2 \operatorname{Re} \{B_n\}} \right] \,, \tag{14}$$

where $\operatorname{Im} \{A_n\}$ and $\operatorname{Re} \{B_n\}$ represent the imaginary and real parts of A_n and B_n, respectively. This inequality determines the minimum frequency for which the matrix is strictly diagonally dominant. For frequencies above this limit, iterative methods should perform better. Although this result is only a sufficient condition, it sheds some light on the poor performance of iterative methods for downward continuation, particularly regarding their low-frequency behavior. The limit of equation (14) depends on the ratio

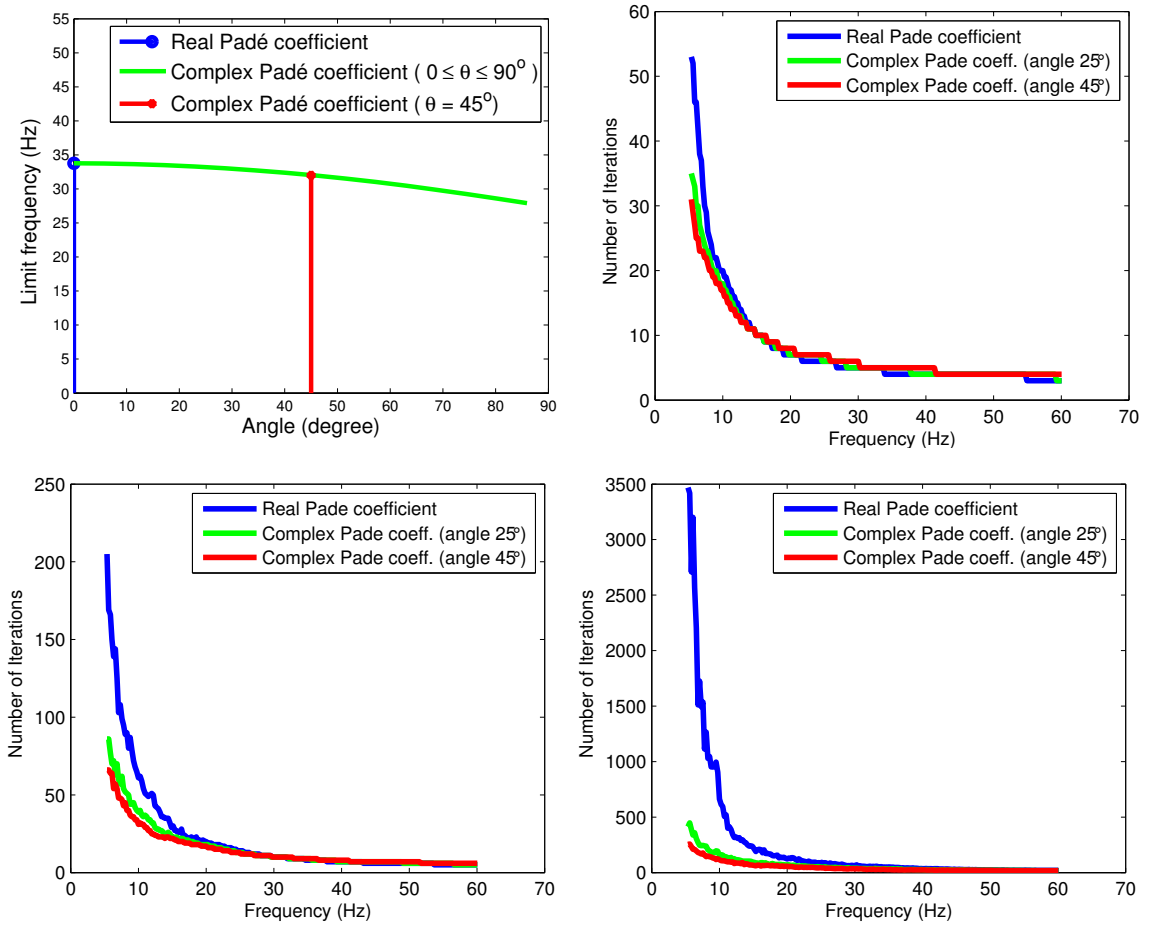

Figure 2: Top left: Limit frequency, f_L, versus rotation angle in degrees for $c/\Delta x = 150$ Hz. Top right: number of iterations of BICGSTAB for each frequency for $c/\Delta x = 75$ Hz. Bottom left: iterations versus frequency for $c/\Delta x = 150$ Hz, Bottom right: iterations versus frequency for $c/\Delta x = 450$ Hz.

between propagation velocity and grid interval. When this ratio increases, i.e., for small grid sizes, the convergence of iterative methods is slower.

Another consequence from equation (14) regards the difference between the real and complex Padé expansions. In the real version, $\mathrm{Im}\{A_n\} = \mathrm{Im}\{B_n\} = 0$, and equation (14) reduces to

$$\omega \geq \omega_L = \frac{2c}{\Delta x}\sqrt{2B_n}\,. \tag{15}$$

This implies that using the complex Padé expansion will improve convergence, and it will improve the more the larger $\mathrm{Im}\{A_n\}$ is relative to $\mathrm{Re}\{B_n\}$. This fact is corroborated in the top left part of Figure 2, which shows how the limit frequency $f_L = \omega_L/2\pi$ depends on the branch cut rotation angle θ, for $N = 1$ and $c/\Delta x = 150$ Hz (e.g., $\Delta x = 10$ m and $c = 1500$ m/s).

To evaluate how the actual behavior of BICGSTAB relates to inequality (14), the remaining graphs in Figure 2 display the number of iterations as a function of frequency for increasing values of the ratio $c/\Delta x$ for real and for complex Padé coefficients with $\theta = 25°$ and $\theta = 45°$. These graphs shows that the complex Padé approximation markedly improves the convergence of BICGSTAB, particularly for large values of the ratio $c/\Delta x$. Increasing the branch-cut rotation angle from $05°$ (real approximation) to $25°$ strongly impacts the convergence, while the increase from $25°$ to $45°$, though further reducing the number of iterations, has a less pronounced effect.

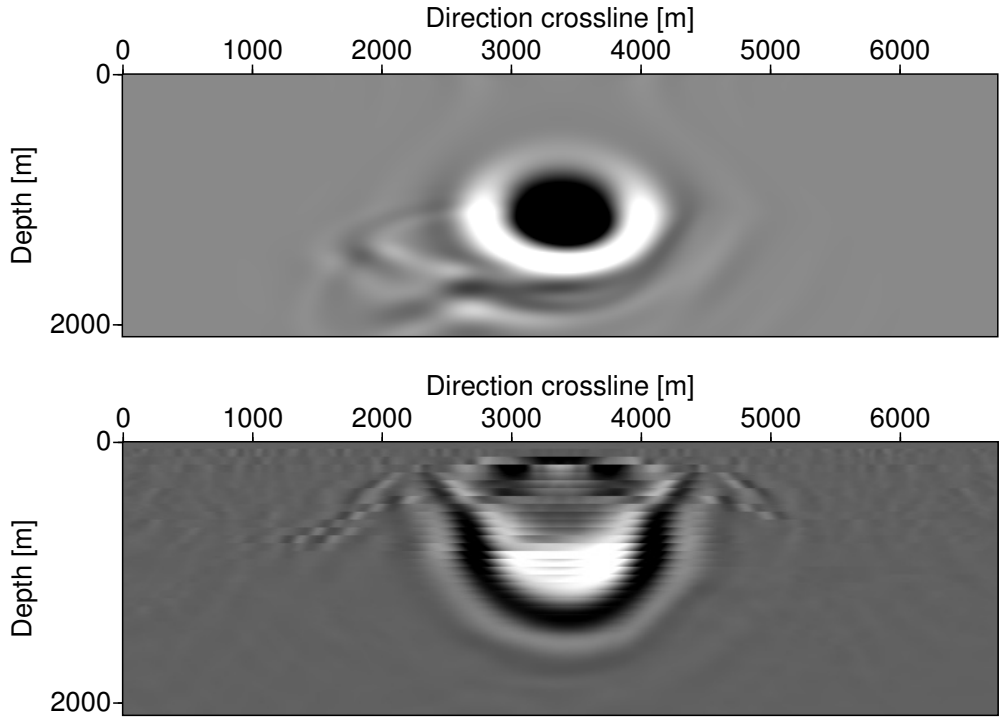

Figure 3: Sections of 3D FD migration impulse response in the plane $x = 1500$ m. Top: BICGSTAB, $\theta = 25^\circ$. Bottom: 2way splitting plus Li correction

NUMERICAL RESULTS

After the evaluation of the performance, we are interested in the quality of the resulting wavefield. As a first test, we computed the impulse response of the 3D FD migration operator in a homogeneous medium with and without inline and crossline splitting. The propagation velocity is 1500 m/s. We used a single term of the complex Padé expansion with rotation angle $\theta = 45^\circ$ to approximate the migration operator. The impulse response without splitting was computed using BICGSTAB. The grid dimensions are $676\times676\times210$ along coordinates x, y, and z, respectively. The grid spacing is uniform and equal to 10 m. The volume injection source is located one grid spacing below the surface in the center of the grid. The source signature is a Ricker pulse with peak frequency 25 Hz, and the sample rate is 8 ms.

The resulting cross-sections of the 3D impulse responses at depth $z = 1050$ m are depicted in Figure 1. The red circle in these images indicate the correct position. The improvement of the impulse response computed by BICGSTAB over the one computed using operator splitting is evident. The correct positioning and the absence of numerical anisotropy indicate the good performance of the BICGSTAB algorithm.

Next, we use the SEG/EAGE salt model to evaluate the performance of 3D migration without splitting in a complex inhomogeneous medium. The strong lateral variations between the sediments and the salt body and the high velocity of the salt are a challenge for the iterative methods. To filter out the spike velocity contrast used to simulate reflectors in the original model we applied median filter of size $7\times7\times7$. The grid size, source position and source pulse are the same as for the homogeneous example. We computed FD impulses responses using BICGSTAB with a single term of the complex Padé expansion $\theta = 25^\circ$ and $\theta = 45^\circ$. These impulse responses are compared to FD impulse response computed using the direct solver MUMPS using three terms of the complex Padé response and $\theta = 45^\circ$. Finally, we compare them to the 3D impulse response of reverse time migration.

For a single term of the complex Padé expansion the MUMPS direct solver is two times slower than BICGSTAB. On the other hand, for higher order complex Padé terms, the performance of BICGSTAB degrades rapidly, and increasing the angle θ does not help in this case. This agrees with the observation of Nichols (1991) that iterative methods have more difficult modeling the wavefield propagating at higher dip

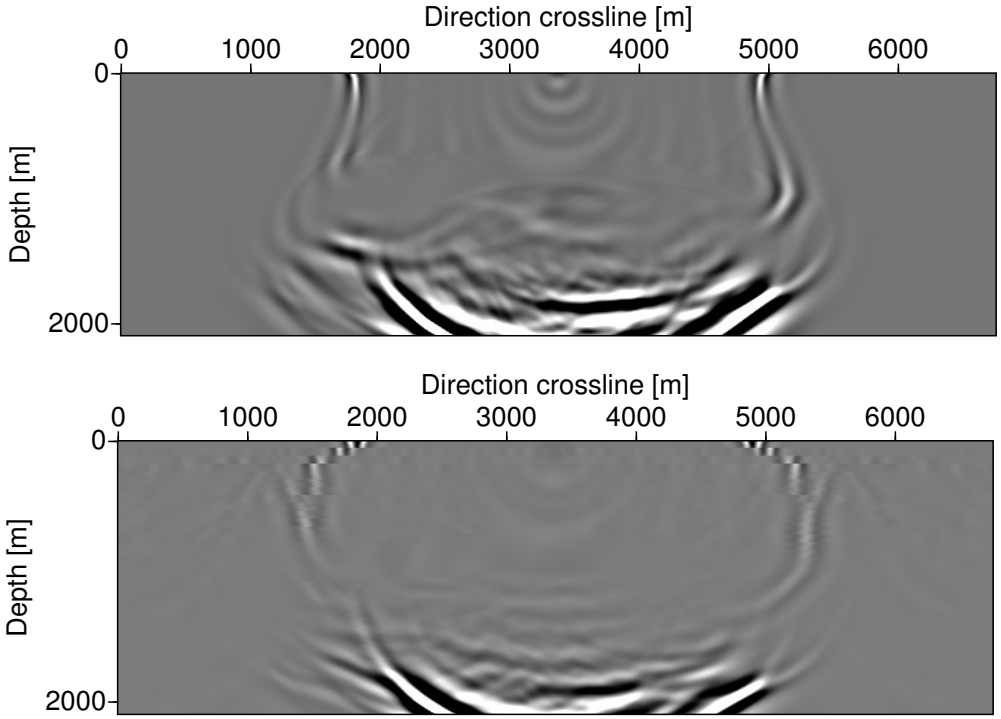

Figure 4: Sections of 3D FD migration impulse response in the plane $x = 3380$ m. Top: BICGSTAB, $\theta = 25°$. Bottom: 2way splitting plus Li correction

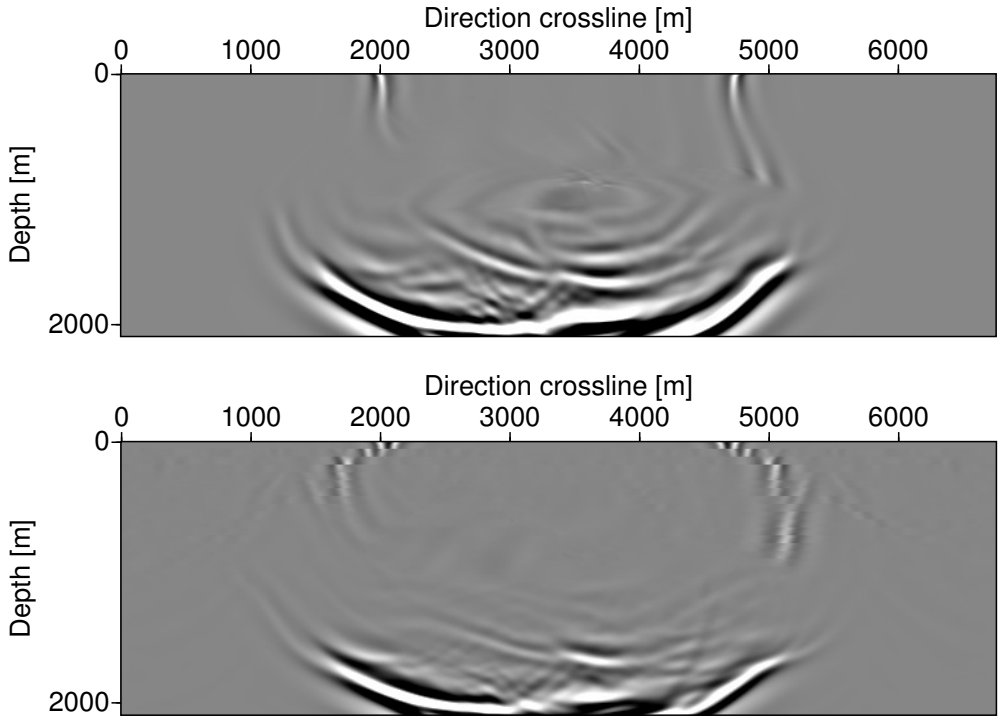

Figure 5: Sections of 3D FD migration impulse response in the plane $x = 4160$ m. Top: BICGSTAB, $\theta = 25°$. Bottom: 2way splitting plus Li correction

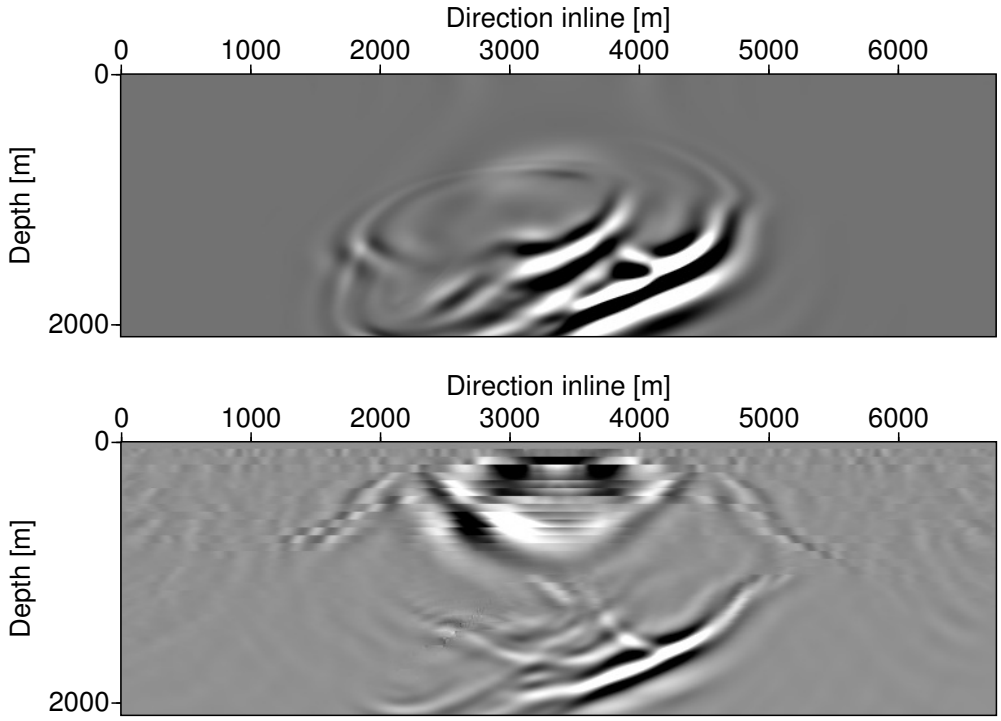

Figure 6: Sections of 3D FD migration impulse response in the plane $y = 1500$ m. Top: BICGSTAB, $\theta = 25°$. Bottom: 2way splitting plus Li correction

angles. Only the direct solver was able to compute the impulse response with three Padé terms.

Figs. 3, 4 and 5 show three vertical sections of the impulse response along the crossline direction (y) at inline positions $x = 1500$ m, $x = 3880$ m, and $x = 4160$ m, computed with 3D FD migration using BICGSTAB with a single complex Padé term with $\theta = 25°$ and compare them to corresponding results obtained with two-way splitting in the inline and crossline directions plus additional Li corrections with 10 reference velocities at every 6 depth steps. When comparing both parts of Figure 3, four important observations can be made: (a) The images differ significantly in the upper part. This is due to different artifacts, since because of its intrinsic restrictions, FD migration cannot properly image events dipping above $45°$. We see that the artifacts of the splitting+correction approach are much stronger than those of the iterative solver. (b) The positioning of the migrated lower dipping events is very similar in both parts of Fig 3. The apparent change in polarity is due to a slightly different positioning of the event in the inline (x) direction. (c) Apart from the different artifacts in the upper part, the Li correction causes a discontinuous behavior of the wavefront. (d) The events of lower amplitude ahead of the main wavefront are not reproduced by the splitting+correction approach. The differences in Figs. 4 and 5 are less pronounced, but of similar nature.

Figs. 6, 7 and 8 show the corresponding comparison of vertical sections along the inline direction (x), at crossline positions $y = 1500$ m, $y = 3880$ m, and $y = 4160$ m, Again, note the good match among corresponding events dipping up to around 45° degrees, but the differences for higher propagation angles, amplitudes, and migration artifacts. As with the crossline sections, the strongest artifacts can be observed in the section at $y = 1500$ m. From these examples, we conclude that the splitting+correction approach has a number of drawbacks that can be avoided by the iterative solver at the expense of a higher computational cost.

In the next set of figures, we compare the BICGSTAB results to those of the direct MUMPS solver. Figs. 9, 10, 11, and 12 show horizontal sections of the 3D impulse response computed by four different methods: FD migration solved using BICGSTAB for one complex Padé term with $\theta = 25°$ (top left) and

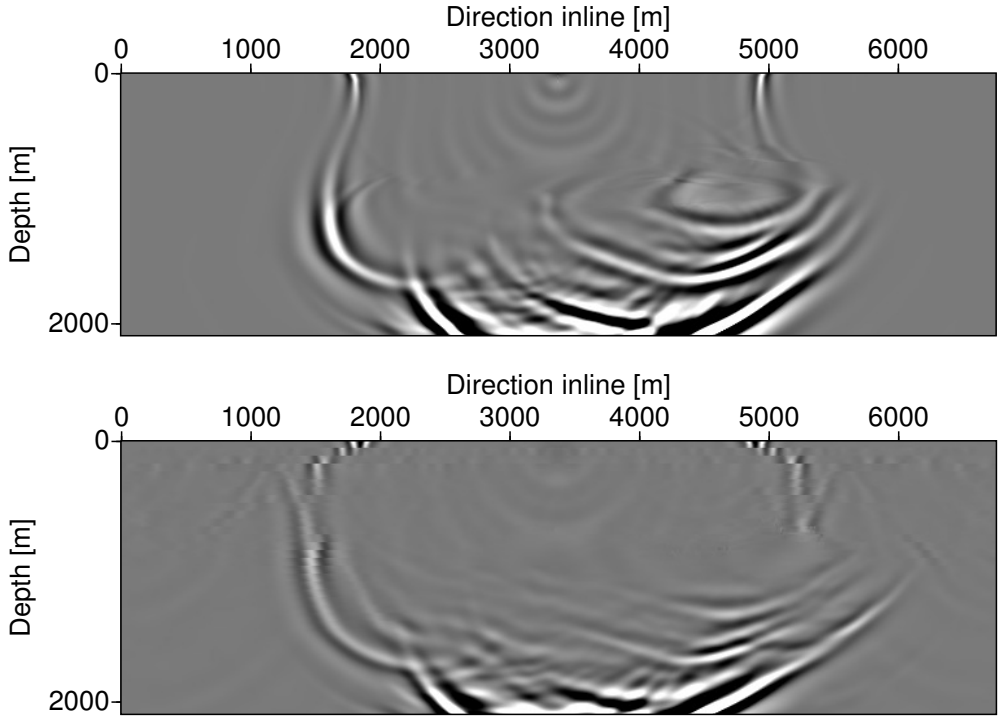

Figure 7: Sections of 3D FD migration impulse response in the plane $y = 3380$ m. Top: BICGSTAB, $\theta = 25°$. Bottom: 2way splitting plus Li correction

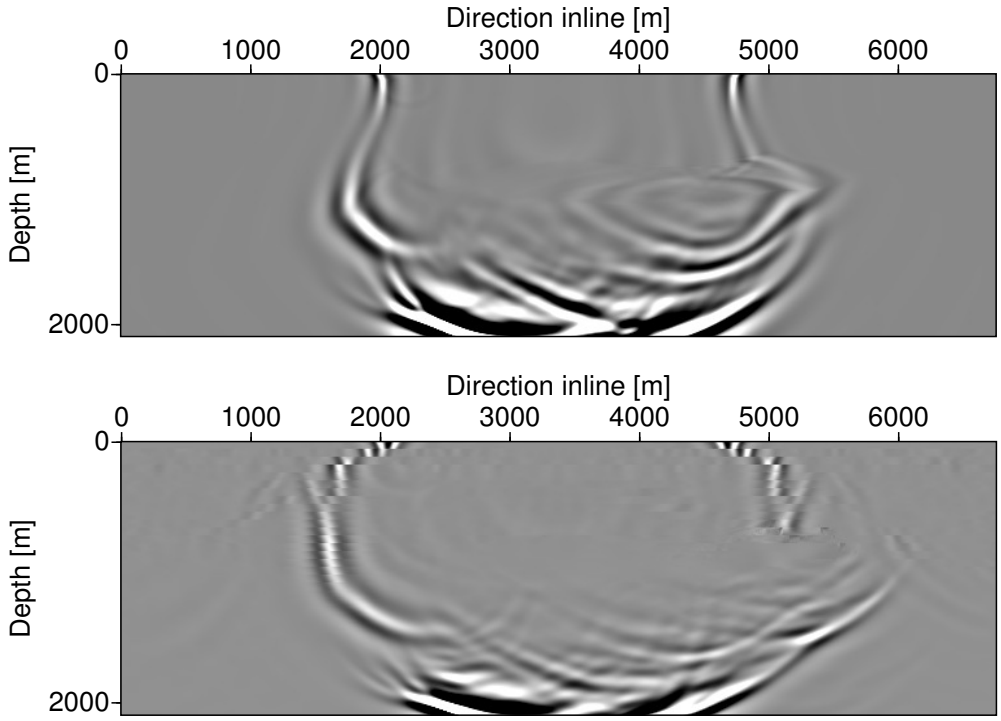

Figure 8: Sections of 3D FD migration impulse response in the plane $y = 4160$ m. Top: BICGSTAB, $\theta = 25°$. Bottom: 2way splitting plus Li correction

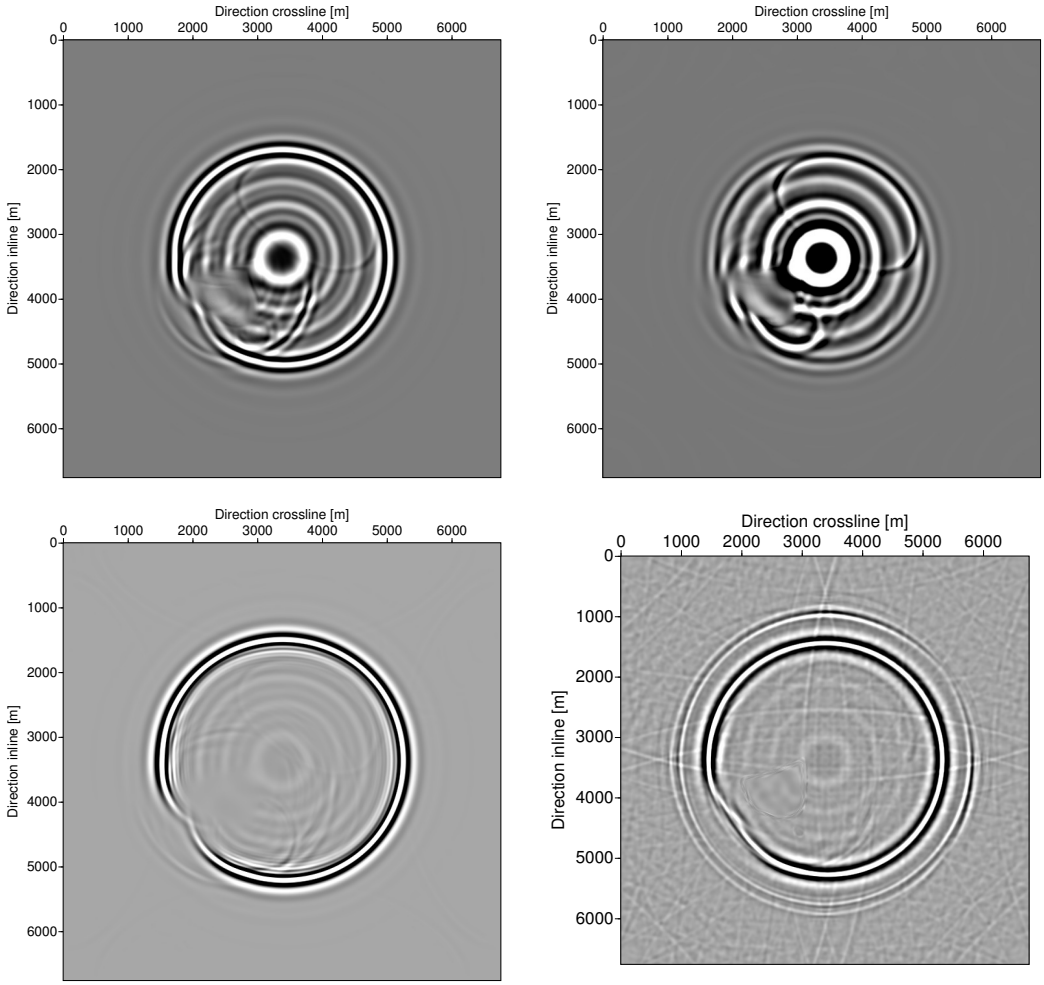

Figure 9: Sections of 3D impulse responses in the plane $z = 550$ m. Top left: FD migration computed using BICGSTAB, $\theta = 25°$. Top right: FD migration using BICGSTAB $\theta = 45°$. Bottom left: FD migration computed using direct solver MUMPS, $\theta = 45°$. Bottom right: 2way splitting plus Li correction.

$\theta = 45°$ (top right), FD migration using the MUMPS direct solver for three complex Padé terms with $\theta = 45°$ (bottom left) and the reference solution using three $\theta = 45°$ Padé terms with two-way splitting plus Li correction (bottom right).

The section in Figure 9 is at a very shallow level. It mainly illustrates the differences of the different methods regarding the migration artifacts at high propagation angles. The main wavefront in the three-term MUMPS and splitting+correction solutions is slightly better positioned because of the 45° dip limit of the single-term complex Padé expansion used in BICGSTAB.

Figs. 10, 11, and 12 demonstrate the quality of the one-term BICGSTAB solution. The impulse response is much closer to the theoretically better three-term MUMPS result than the one obtained with two-way splitting plus Li correction. Note again the noncausal artifacts of the latter approach ahead of the wavefront. From these figures, we see that the difference among corresponding events between MUMPS and BICGSTAB diminishes with depth, indicating that the iterative solver is sufficiently accurate for propagation directions below 45°. The differences between the $25°$ and $45°$ rotation angle BICGSTAB solutions are rather subtle.

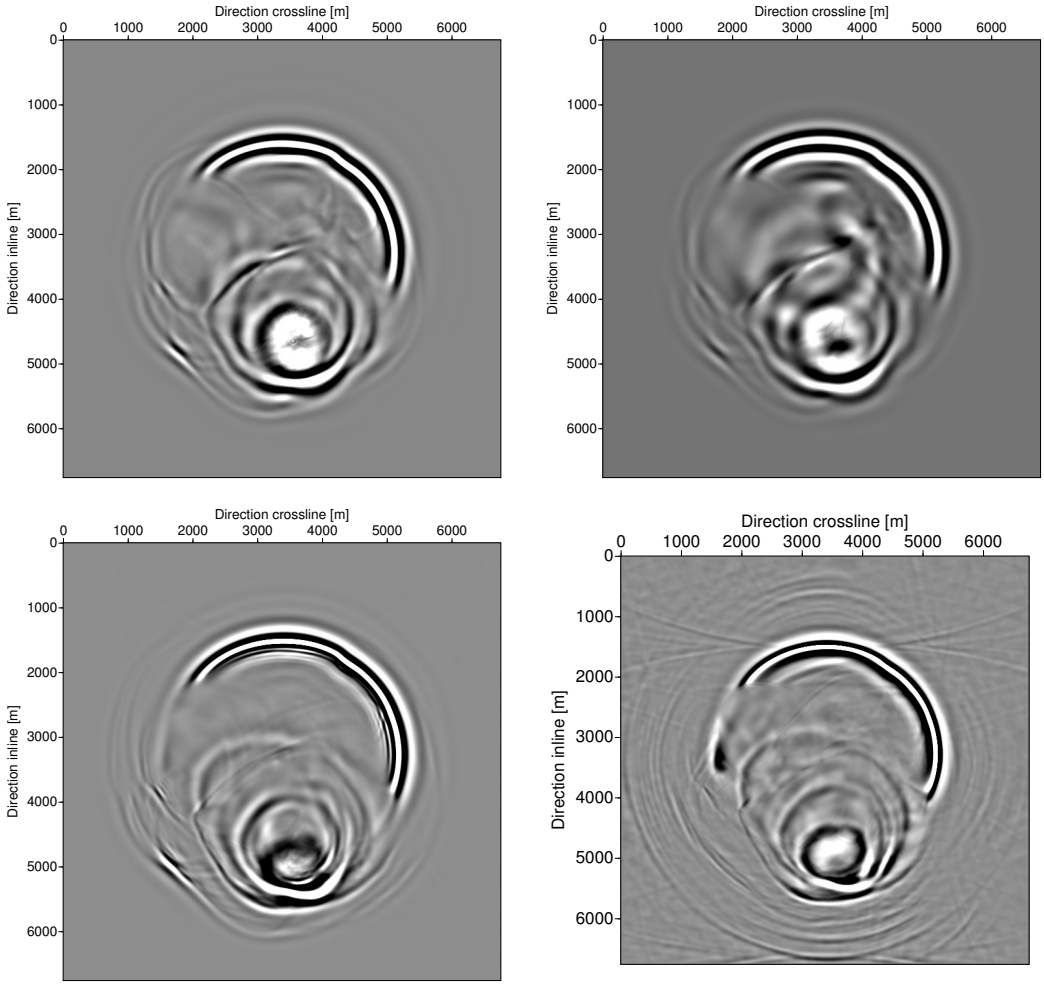

Figure 10: Sections of 3D impulse responses in the plane $z = 1050$ m. Top left: FD migration computed using BICGSTAB, $\theta = 25°$. Top right: FD migration using BICGSTAB $\theta = 45°$. Bottom left: FD migration computed using direct solver MUMPS, $\theta = 45°$. Bottom right: 2way splitting plus Li correction.

COMPUTATION TIME

When comparing splitting and nonsplitting solvers of 3D FD wave-equation migration, a word on computation time is indispensable. In our implementation, the BICGSTAB solution for one-term complex Padé FD migration took half as much time as the corresponding MUMPS solution and about five times as much as a corresponding solution based on two-way splitting plus Li correction with 10 reference velocities. The three-term MUMPS solution consumed three times more time than the one-term solution. The three-term BICGSTAB solution converges in a much longer time than the three terms with MUMPS solution. In other words, of the methods used in the above figures, BICGSTAB with a single term is intermediate between the six times more time-consuming three-term MUMPS solution and the five times less time-consuming splitting+correction approach.

CONCLUSIONS

The performance of iterative methods for 3D downward continuation without splitting along inline and crossline directions improves markedly using the complex Padé approximation. Numerical experiments using the SEG/EAGE salt model show the feasibility of FD migration without splitting for such a complex

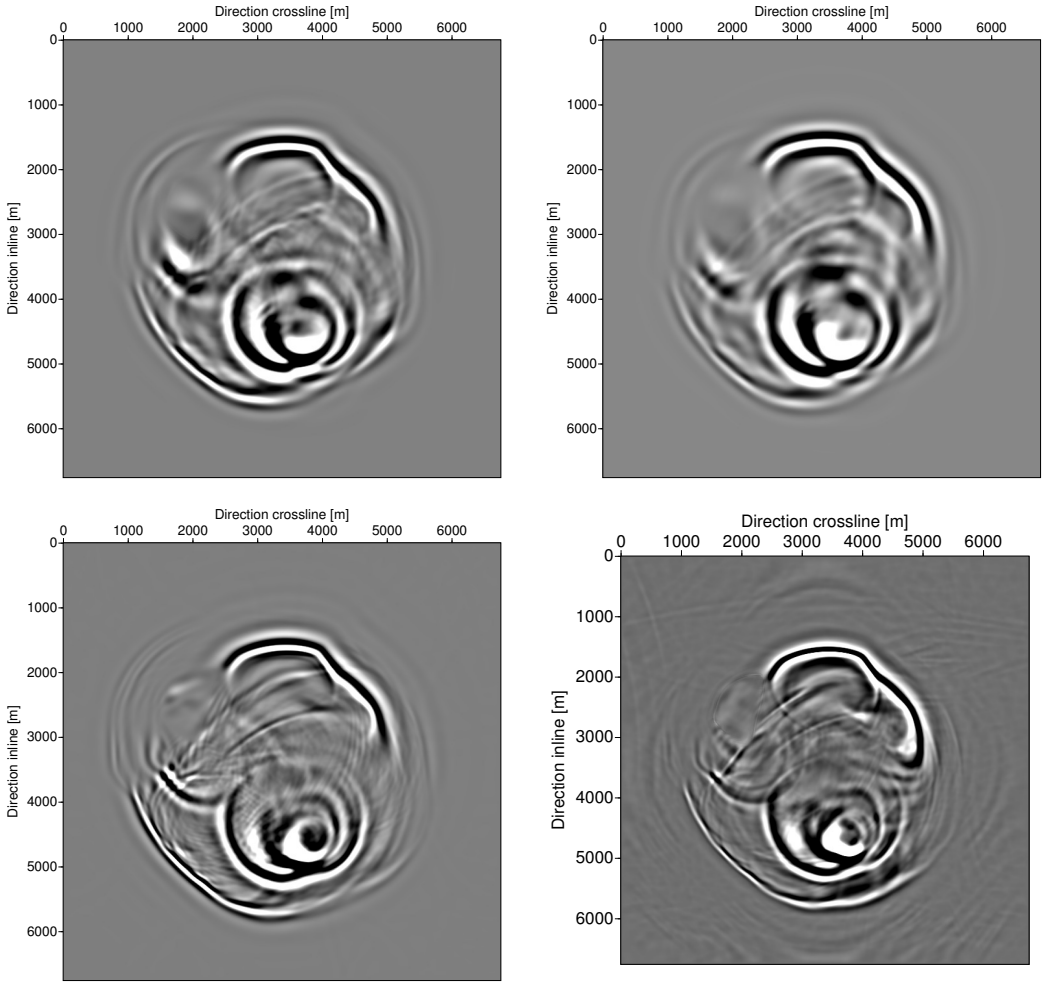

Figure 11: Sections of 3D impulse responses in the plane $z = 1350$ m. Top left: FD migration computed using BICGSTAB, $\theta = 25°$. Top right: FD migration using BICGSTAB $\theta = 45°$. Bottom left: FD migration computed using direct solver MUMPS, $\theta = 45°$. Bottom right: 2way splitting plus Li correction.

velocity model, being less than an order of magnitude slower than conventional splitting approaches. We provided a simple analysis indicating that the complex Padé expansion improves the conditioning of FD downward continuation. For a single term in the Padé expansion, BICGSTAB outperforms the massively parallel direct solver by a factor of two. The performance of iterative methods degrades rapidly when more terms of the Padé are used but does not affect MUMPS in the same way. The computation time of MUMPS increases only linearly with the number of terms.

Efficiency of full solutions without operator splitting still remains an issue. Our results indicate that the combination of the complex Padé approximation with additional preconditioning can further reduce the computational cost, possibly bringing non-splitting solutions within the range of more sophisticated corrections to splitting. Moreover, since low frequency is the main problem affecting the performance of BICGSTAB, multigrid techniques can certainly help to speed up these methods.

ACKNOWLEDGMENTS

We acknowledge the financial support from PETROBRAS and the sponsors of the Wave Inversion Technology (WIT) consortium, as well as the Brazilian Agencies CAPES, CNPq, and FINEP.

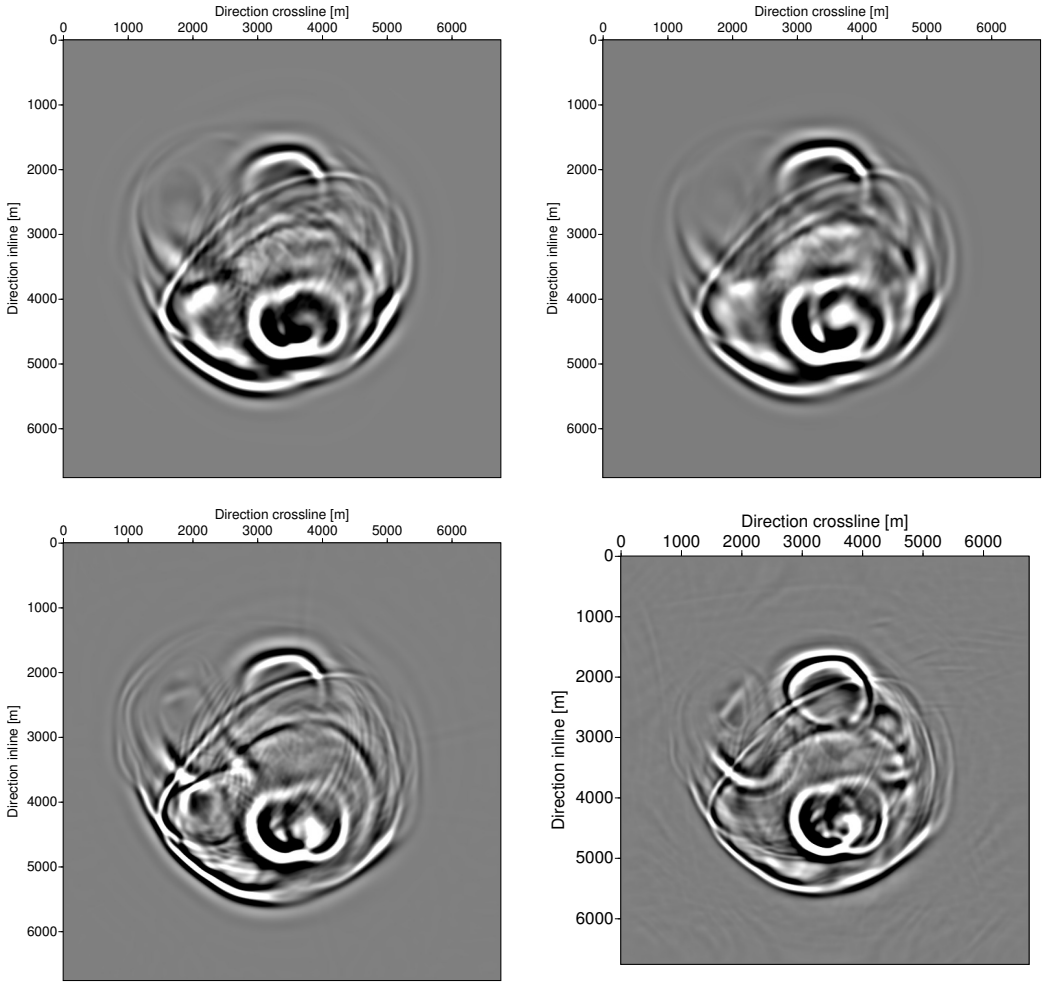

Figure 12: Sections of 3D impulse responses in the plane $z = 1550$ m. Top left: FD migration computed using BICGSTAB, $\theta = 25°$. Top right: FD migration using BICGSTAB $\theta = 45°$. Bottom left: FD migration computed using direct solver MUMPS, $\theta = 45°$. Bottom right: 2way splitting plus Li correction.

REFERENCES

Amazonas, D., Costa, J. C., Schleicher, J., and Pestana, R. (2007). Wide-angle FD and FFD migration using complex padé approximations. *Geophysics*, 72(6):S215–S220.

Amestoy, P. R., Duff, I. S., Koster, J., and L'Excellent, J.-Y. (2001). A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM Journal of Matrix Analysis and Applications*, 23(1):15–41.

Amestoy, P. R., Guermouche, A., L'Excellent, J.-Y., and Pralet, S. (2006). Hybrid scheduling for the parallel solution of linear systems. *Parallel Computing*, 32(2):136–141.

Bamberger, A., Engquist, L. H., and Joly, P. (1988). Higher order paraxial wave equation approximations in heterogeneous media. *J. Appl. Math.*, 48:129–154.

Brown, C. L. (1983). Application of operator separation in reflection seismology. *Geophysics*, 48(3):288–294.

Claerbout, J. F. (1985). *Imaging the Earth's Interior*. Blackwell Sci. Pub.

Cole, S. (1989). Iterative methods for 3-D finite-difference migration and modeling. Technical report, Stanford University.

Iserles, A. (1996). *A first course in the numerical analysis of differential equations*. Cambridge University Press.

Li, Z. (1991). Compensating finite-difference migration by multiway splitting. *Geophysics*, 56(10):1650–1660.

Millinazzo, F. A., Zala, C. A., and Brooke, G. H. (1997). Square-root aproximations for parabolic equation algorithms. *J. Acoust. Soc. Am.*, 101(2):760–766.

Nichols, D. (1991). 3-D depth migration by a predictor corrector method. Technical report, Stanford University.

Ristow, D. and Rühl, T. D. (1997). 3D finite-difference erros in 3-D migration and modeling. *Geophysics*, 62(2):554–567.

van der Vorst, H. A. (1992). Bi-CGSTAB: a fast and smoothly converging variant of BI-CG for nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, 13(2):631–644.