

# Acoustic Staggered Grid Modeling in IWAVE

*William. W. Symes*

## ABSTRACT

IWAVE is a framework for time-domain regular grid finite difference and finite element methods. The IWAVE package includes source code for infrastructure component, and implementations of several wave physics modeling categories. This paper presents two sets of examples using IWAVE acoustic staggered grid modeling. The first set illustrates the effectiveness of a simple version of Perfectly Matched Layer absorbing boundary conditions. The second set reproduce illustrations from a recent paper on error propagation for heterogeneous medium simulation using finite differences, and demonstrate the interface error effect which renders all FD methods effectively first-order accurate. The source code for these examples is packaged with the paper source, and supports the user in duplicating the results presented here and using IWAVE in other settings.

## INTRODUCTION

Domain-specific simulation such as seismic modeling begs for software re-use via modular design. All applications of this type have the same structure: static fields are initialized, dynamic fields updated, output extracted. A modular approach to code architecture is implicit in this structure, and further specialization leads to even more opportunity for code re-use via modular design.

IWAVE is open source software for finite difference or finite element time-domain simulation on regular rectangular grids, written exclusively in the C/C++. IWAVE is built around a core framework: that is, a collection of separate software packages which together provide essential services upon which applications may be built. These service components completely define the interfaces to which additional code must be written to formulate a complete application. The core framework defines

- parameter-driven job control;
- grid generation and memory allocation in 1D, 2D, and 3D space;
- serial, loop-parallel, and task-parallel execution models, scaling to thousands of threads;
- arbitrary source and receiver locations, and flexible source specification including simultaneous source modeling (random, plane-wave,...)

- standard input and output data formats (SEG-Y, RSF)
- predefined support for linearized (Born) modeling and adjoint linearized (RTM) modeling, both first and second order;
- uniform interface to optimization and linear algebra for creation of inversion applications via the Rice Vector Library (“RVL”) (Padula et al., 2009; Symes et al., 2011).

Symes (2014) describes the design principles underlying the IWAVE core framework, and illustrates the construction of a complete acoustic modeling application using centered finite differences for the second order acoustic constant density wave equation.

The primary purpose of this short paper is to illustrate synthetic seismogram generation using another finite difference scheme implemented in IWAVE, the staggered grid approximation to variable-density velocity-pressure acoustodynamics (Virieux, 1984). Exactly the same framework supports this application as was described in Symes (2014); as explained there, only two data structures and six principal functions need be defined to implement this (or any) finite difference method in IWAVE.

The examples illustrate two aspects of finite difference modeling. The IWAVE staggered grid implementation includes a version of PML absorbing boundary conditions (Hu et al., 2007), permitting accurate finite grid approximation of wave propagation in a full- or half-space. The first set of examples demonstrates the effectiveness of these very simple PML conditions. The second set reproduce the examples presented in Symes and Vdovina (2009), and illustrate a fundamental limitation in the use of straightforward finite-difference methods for modeling waves in heterogeneous media.

IWAVE was used in a quality control role in the SEAM Phase I project - see Fehler and Keliher (2011) for an account, including discussion of the many difficulties of large scale numerical simulation of seismograms.

The internal details of IWAVE are not discussed here, except insofar as is necessary to explain the use of the main commands. As mentioned above, Symes (2014) overviews the design of IWAVE and the main features of its internal structure, and defines the elements necessary to compile a new IWAVE application. Symes et al. (2011) briefly describe the IWAVE/RVL mechanisms for coupling modeling with optimization packages to produce inversion applications.

The paper begins with a brief review of the system of partial differential equations solved (approximately) by IWAVE’s acoustic application, and the choice of finite difference method. The next section evaluates the effectiveness of the PML absorbing boundary conditions included in the IWAVE staggered grid acoustic application. The following section presents the examples of Symes and Vdovina (2009), along with some additional examples based on the same distribution of mechanical parameters which shed light on the impact of finite difference order on solution accuracy. Instructions

follow for recreating these examples, and for using them as starting points for further modeling exercises. The paper ends with a brief discussion of the prospects for improvements in performance and accuracy in FD technology, and the evolutionary advantages flowing from the modular, or object, orientation of IWAVE. Two appendices describe the job parameters used in the examples, and download and install instructions.

## ACOUSTODYNAMICS

The pressure-velocity form of acoustodynamics consists of two coupled first-order partial differential equations:

$$\rho \frac{\partial \mathbf{v}}{\partial t} = -\nabla p \quad (1)$$

$$\frac{1}{\kappa} \frac{\partial p}{\partial t} = -\nabla \cdot \mathbf{v} + g \quad (2)$$

In these equations,  $p(\mathbf{x}, t)$  is the pressure (excess, relative to an ambient equilibrium pressure),  $\mathbf{v}(\mathbf{x}, t)$  is the particle velocity,  $\rho(\mathbf{x})$  and  $\kappa(\mathbf{x})$  are the density and bulk modulus respectively. Bold-faced symbols denote vectors; the above formulation applies in 1, 2, or 3D.

The inhomogeneous term  $g$  represents externally supplied energy (a “source”), via a defect in the acoustic constitutive relation. A typical example is the *isotropic point source*

$$g(\mathbf{x}, t) = w(t)\delta(\mathbf{x} - \mathbf{x}_s)$$

at source location  $\mathbf{x}_s$ .

Virieux (1984) introduced finite difference methods based on this formulation of acoustodynamics to the active source seismic community. Virieux (1986) extended the technique to elastodynamics, and Levander (1988) demonstrated the use of higher (than second) order difference formulas and the consequent improvement in dispersion error. Many further developments are described in the review paper Moczo et al. (2006). IWAVE’s acoustic application uses the principles introduced by these authors to offer a suite of finite difference schemes, all second order in time and of various orders of accuracy in space.

The bulk modulus and buoyancy (reciprocal density) are the natural parameters whose grid samplings appear in the difference formulae. These are the parameters displayed in the figures below, rather than, say, velocity and density, which might seem more natural.

## PML EFFECTIVENESS

The IWAVE acoustic staggered grid scheme implements the Perfectly Matched Layer (PML) approach to absorbing boundary conditions, in one of the simpler of its many

guises (a split field approach - (Hu et al., 2007)). After some manipulation, the acoustic PML system for the physical velocity  $\mathbf{v}$  and an artificial vector pressure  $\mathbf{p}$  takes the form

$$\begin{aligned} \rho \left( \frac{\partial v_k}{\partial t} + \eta_k(x_k)v_k \right) &= -\frac{\partial p_k}{\partial x_k}, \\ \frac{1}{\kappa} \left( \frac{\partial p_k}{\partial t} + \eta_k(x_k)p_k \right) &= -\nabla \cdot \mathbf{v} + g \end{aligned} \quad (3)$$

in which the  $k$ th component of the attenuation profile vector  $\eta$  depends only on  $x_k$ , and can be stored as a collection of 1D objects. Ordinary acoustic wave propagation takes place where  $\eta = \mathbf{0}$ , and if the components of the vector pressure  $\mathbf{p}$  are all the same in this zone, then they remain the same there, and any one of them may be regarded as the same as the physical pressure field. Outside of the physical domain, where waves are to be attenuated,  $\eta$  should be positive; at the boundary of the physical domain, it should vanish to positive order. We elected to make  $\eta$  cubic in distance to the boundary: for a PML layer of width  $L_{k,r}$ , beginning at  $x_k = x_{k,r}$  along the  $k$ th coordinate axes,

$$\eta(x_k) = \eta_0 \left( \frac{x_k - x_{k,r}}{L_{k,r}} \right)^3$$

Thus there are four PML boundary layer thicknesses in 2D, six in 3D, one for each side of the simulation cube. The IWAVE convention imposes pressure-free boundary conditions on the exterior boundary of the PML domain. Thus  $L = 0$  signifies a free surface boundary face. Any face of the boundary may be assigned a zero-pressure condition ( $L = 0$ ) or a PML zone of any width ( $L > 0$ ).

Many implementations of PML, especially for elasticity, confine the extra PML fields (in this case, the extra pressure variables) to explicitly constructed zones around the boundary, and use the standard physical system in the domain interior. We judged that for acoustics little would be lost in either memory or efficiency, and much code bloat avoided, if we were to solve the system (3) in the entire domain.

Considerable experience and some theory (Hu et al., 2007; Moczo et al., 2006) suggest that the system 3 will effectively absorb waves that impinge on the boundary, emulating free space in the exterior of the domain, if the PML zones outside the physical domain in which  $\eta$  are roughly a half-wavelength wide, and  $\eta_0 = 0$ .

A simple 2D example illustrates the performance of this type of PML. The physical domain is a 1.8 x 7.6 km; the same domain is used in the experiments reported in the next section. A point source is placed at  $z=40$  m,  $x = 3.3$  km, with a Gaussian derivative time dependence with peak amplitude at about 5 Hz, and significant energy at 3 Hz but little below. The acoustic velocity is 1.5 km/s throughout the domain, so the effective maximum wavelength is roughly 500 m. The density is also constant, at 1 g/cm<sup>3</sup>. A snapshot of the wavefield at 1.2 s after source onset (Figure 1), before the wave has reached the boundary of the domain, shows the expected circular wavefront. At 4.0 s, a simulation with zero-pressure boundary conditions on all sides

of the physical domain produces the expected reflections, Figure 2. With PML zones of 250 m on the bottom and sides of the domain, so that only the top is a zero-pressure surface, and  $\eta_0 = 1$ , the wave and its free-surface ghost both appear to leave the domain (Figure 3, plotted on the same grey scale). The maximum amplitude visible in Figure 2 is roughly  $7.1 \times 10^{-2}$ , whereas the maximum amplitude in Figure 3 is  $7.0 \times 10^{-5}$ . The actual reflection coefficient is likely less than  $10^{-3}$ , as the 2D free space field does not have a lacuna behind the wavefront, but decays smoothly, so the low end of the wavelet spectrum remains.

It is not possible to decrease the PML layer thickness much beyond the nominal longest half-wavelength and enjoy such small reflections. Figure 4 shows the field at 4.0 s with PML zones of width 100 m on bottom and sides, and an apparently optimal choice of  $\eta_0$ . The maximum amplitude is  $2.3 \times 10^{-4}$ , and a reflected wave is clearly visible at the same grey scale.

## ALL FD SCHEMES ARE FIRST ORDER IN HETEROGENEOUS MEDIA

The bulk modulus and buoyancy models depicted in Figures 5 and 6 embed an anticline or dome in an otherwise undisturbed package of layers. These figures display sampled versions of the models with  $\Delta x = \Delta z = 5$  m; the model fields are actually given analytically, and can be sampled at any spatial rate. The IWAVE utility *sfstandardmodel* (in the Madagascar `bin` directory) builds this example and a number of others that can be sampled arbitrarily for grid refinement studies. See its self-doc for usage instructions.

Symes and Vdovina (2009) use the model depicted in Figures 5 and 6 to illustrate the *interface error* phenomenon: the tendency, first reported by Brown (1984), of all finite difference schemes for wave propagation to exhibit first order error, regardless of formal order, for models with material parameter discontinuities. Figure 7 exhibits a shot gather, computed with a (2,4) (= 2nd order in time, 4th order in space) staggered grid scheme,  $\Delta x = \Delta z = 5$  m (more than 20 gridpoints per wavelength at the wavelength corresponding to the highest frequency, 12 Hz, with significant energy, and the smallest  $v_p = 1.5$  km/s) and an appropriate near-optimal time step, acquisition geometry as described in caption. The same gather computed at different spatial sample rates seem identical, at first glance, however in fact the sample rate has a considerable effect. Figures 8 and 9 compare traces computed from models sampled at four different spatial rates (20 m to 2.5 m), with proportional time steps. The scheme used is formally 2nd order convergent like the original 2nd order scheme suggested by Virieux (1984), but has better dispersion suppression due to the use of 4th order spatial derivative approximation. Nonetheless, the figures clearly show the first order error, in the form of a grid-dependent time shift, predicted by Brown (1984).

Generally, even higher order approximation of spatial derivatives yields less dis-

persive propagation error, which dominates the finite difference error for smoothly varying material models. For discontinuous models, the dispersive component of error is still improved by use of a higher order spatial derivative approximation, but the first order interface error eventually dominates as the grids are refined. Figure 10 shows the same shot gather as displayed earlier, with the same spatial and temporal sampling and acquisition geometry, but computed via the (2,8) (8th order in space) scheme. The two gather figures are difficult to distinguish. The trace details (Figures 11, 12) show clearly that while the coarse grid simulation is more accurate than the (2,4) result, but the convergence rate stalls out to 1st order as the grid is refined, and for fine grids the (2,4) and (2,8) schemes produce very similar results: dispersion error has been suppressed, and what remains is due to the presence of model discontinuities.

See Symes and Vdovina (2009) for more examples, analysis, and discussion, also Fehler and Keliher (2011) for an account of consequences for quality control in large-scale simulation.

Note that the finest (2.5 m) grid consists of roughly 10 million gridpoints. Consequently the modeling runs collectively take a considerable time, from a minutes to a substantial fraction of an hour depending on platform, on a single thread. This example is computationally large enough that parallelism via domain decomposition is worthwhile. IWAVE is designed from the ground up to support parallel computation; a companion report will demonstrate parallel use of IWAVE.

## CREATING THE EXAMPLES - RUNNING IWAVE APPLICATIONS

IWAVE builds with SConstruct (<http://www.scons.org>), either as an independent package or as part of Madagascar (Fomel, 2009). See the Madagascar web site

[http://www.ahay.org/wiki/Main\\_Page](http://www.ahay.org/wiki/Main_Page)

for download and install instructions. Source for IWAVE and other TRIP software reside in the `trip` subdirectory of the top-level Madagascar source directory. A `README` file describes how to install TRIP software independently of the rest of Madagascar, which is useful to configuring TRIP differently from other parts of the package (for example, with MPI support).

The IWAVE acoustic staggered grid modeling command is `sfasg` for the Madagascar build, stored in the Madagascar `$$RSFROOT/bin` directory, or

`$$RSFSRC/trip/iwave/asg/main/asg.x`

for the standalone build. All IWAVE commands self-document: entering the command path prints a usage statement to the terminal, including descriptions of all parameters.

The paper you are currently reading follows the reproducible research pattern described on the Madagascar web site, using Madagascar reproducible research tools. You can find the LaTeX source in the subdirectory `book/trip/asg` of the Madagascar source directory, and the script for building the data in

```
$RSFSRC/book/trip/asg/project/SConstruct
```

This script, together with the self-doc for the acoustic staggered grid command and the remarks in the remainder of this section, should enable you to build your own examples after the pattern used in this project.

IWAVE applications currently expect model data files in the RSF format of Madagascar (Fomel, 2009). Data from other sources will need conversion to this format. An RSF data set consists of two files, an ascii header (grid metadata) file and a flat binary data file. The data set is referenced by the header file name; one of the parameters listed in the header file is the pathname of the binary data file, with key `in`. The header file is small and easily created by hand with an editor, if necessary. Madagascar commands add processing history information to header files, and modify their parameters. By convention, the last value of a parameter (`key=value` pair) appearing in the file is the current value. Many archival data formats make the grid sample values available as a flat binary file - this is true for instance of the gridded models output by GOCAD (<http://www.gocad.org>), for which the `vo` files contain virtually the same information as (so may easily be translated to) RSF header files in ascii form, and the `vodat` files are flat binary files which may be used unaltered as RSF binary files.

IWAVE uses two extensions of the Madagascar RSF standard. The first is the optional inclusion of the `dim` and `gdim` keywords. These permit IWAVE applications to treat an RSF file image as defining a `gdim` dimensional data hypercube divided into `dim` dimensional slices. The second is the axis identification keyword set, `id1`, `id2`, etc.: these supply information on the physical meaning of various axes. For an IWAVE `dim` space-dimensional modeling problem, axes labeled `id1`, ..., `id[dim-1]` are the spatial grid axes. If `gdim`  $\geq$  `dim`, then `id[dim]` labels the time axis, and `id[n]`, `n > dim`, axes other than those of space-time. The IWAVE structure paper (Symes, 2014) explains the use of the additional keywords in more detail.

An example of this construction appears in the script that builds the PML examples above, which are actually frames of movies. The output of the 2D simulations are 3D RSF files (`gdim=3`, `dim=2`) with `id3=2`, that is, the third axis is treated as time. Madagascar applications ignore these keywords: in particular, you can view the 3D RSF simulation output as a movie using `sfgrey` and `xtpen` as usual. The presence of the additional keywords is necessary in order for IWAVE to correctly interpret the data geometry.

This example illustrates another important feature of IWAVE applications: any output data files must exist prior to execution - their data samples are overwritten. The `SConstruct` for this project uses `sfmakevel` to create the movie output files

and `sfput` to add the IWAVE-specific keywords to the headers, before invoking the IWAVE command.

By IWAVE convention, the dimension of the problem is that of the primary model grid. In the acoustic staggered grid application, the primary model grid is that associated with the bulk modulus data. This grid is also the primary grid of the simulation: that is, the space steps used in the finite difference method are precisely those of the bulk modulus data. Thus the choice of simulation grid is made externally to IWAVE.

The IWAVE acoustic application uses specific internal scales - m/ms for velocity, g/cm<sup>3</sup> for density, and corresponding units for other parameters. To ensure that data in other (metric) units are properly scaled during i/o, the RSF header file may specify a value for the `scale` key, equal to the power of 10 by which the data should be multiplied on being read into the application, to convert to the internal scale. For example, if velocities are given in m/s, the header file should include the line `scale = -3`. In forthcoming releases, this device will be deprecated in favor of explicit unit specifications.

The current release is configured to use Seismic Unix (“SU”) (SEG-Y without reel header) format for trace data i/o. Units of length and time are m and ms respectively, consistent with other internal unit choices. Two peculiarities of which the user should be aware: (i) *receiver* coordinates (`gx`, `gy`, and `gelev` keywords) *always* specify trace location, that is, the location at which values are sampled in space-time, and (ii) on input, traces are regarded as point sources, so that each trace multiplies a discrete spatial delta (hence values are scaled by the reciprocal grid cell volume). Both of these design choices stem from the migration (adjoint modeling) and inversion uses of IWAVE, discussed for example in (Symes et al., 2011; Symes, 2014).

Source traces must be modified to conform to this rubric. The `sftowed_array` application relieves the user of the necessity to manually adjust the headers of an SU file containing source traces. It accepts three arguments: (i) an input source file containing `gx`, `gy`, and `gelev` values representing source trace location *relative* to a source center location - the source coordinates of source traces are ignored; (ii) a data file whose `sx`, `sy`, and `selev` values are the source center locations to be used - its receiver coordinates are ignored, and (iii) an output file (name), to which output source traces will be written, each with *source* coordinates equal to those of a data trace, and *receiver* coordinates equal to the sums of the source trace *receiver* coordinates and the data trace *source* coordinates. The result is a collection of source coordinate gathers with the same source coordinates as the data file, but within each gather the same receiver coordinates *relative* to the source coordinates as the source file. Thus the source array is translated to each of the source centers specified in the data file. Because the source file may contain arbitrarily many traces with arbitrary relative locations, any source radiation pattern may be approximated (Santosa and Symes, 2000).

The example scripts in the `project` subdirectory use Madagascar commands to create these prototype trace files.

One of IWAVE’s design criteria is that acquisition geometry parameters should have no *a priori* relation to the computational grid geometry: source and receiver locations may be specified anywhere in Euclidean space.

## DISCUSSION AND CONCLUSION

The rather large and only slowly disappearing error revealed by the examples from Symes and Vdovina (2009) suggests strong limits for the accuracy of regular grid finite difference methods. Finite element methods suffer from the same limitations: accurate solution of acoustodynamic or elastodynamic problems appears to demand interface-fitted meshed (Cohen, 2002), with the attendant increase in code and computational complexity.

The situation may not be so bleak, however. For one special case, namely constant density acoustics, Terentyev and Symes (2009) show that a regular grid finite difference method, derived from a regular grid Galerkin finite element method, has accuracy properties one would expect in homogeneous media (second order convergence, reduction of grid dispersion through higher order space differencing) even for discontinuous models: the interface error effect is attenuated. This type of result actually goes quite far back in computational geophysics (see for example Muir et al. (1992)), though theoretical support has been slower in coming.

Pure regular grid methods cannot take advantage of changes in average velocity across the model, and concomitant changes in wavelength. Coupling of local regular grids is possible, however, and can yield substantial computational efficiency through grid coarsening in higher velocity zones - see Moczo et al. (2006). IWAVE already accommodates multiple grids (in domain decomposition parallelism), and extension to incommensurable multiple grids would be a significant change, but in principle straightforward. The use of logically rectangular but geometrically irregular (“stretched”) grids is completely straightforward, on the other hand.

These and other extensions, both past and future, are eased by the reusability designed into the IWAVE core framework. This design has produced reasonably well-performing and easy-to-use applications, and has proven extensible to new models and schemes. Moreover, as explained by Symes et al. (2011), the object-oriented design of IWAVE dovetails with similarly designed optimization software to support the construction of waveform inversion software. The inversion applications resulting from this marriage inherit the features of IWAVE - parallel execution, high-order stencils, efficient boundary conditions, simple job control - without requiring that these aspects be reworked in the code extensions.

## ACKNOWLEDGEMENTS

IWAVE has been a team effort: the original design of the core framework is due to Igor Terentyev, and Tanya Vdovina, Dong Sun, Marco Enriquez, Xin Wang, Yin Huang, Mario Bencomo, and Muhong Zhou have each made major contributions. Development of IWAVE was supported by the SEG Advanced Modeling (SEAM) project, by the National Science Foundation under awards 0620821 and 0714193, and by the sponsors of The Rice Inversion Project. The IWAVE project owes a great deal to several open source seismic software packages (Seismic Un\*x, SEPlib, Madagascar), debts which we gratefully acknowledge. The author wishes to record his special gratitude to Sergey Fomel, the architect of Madagascar, for his inspiring ideas and his generous and crucial help in the integration of IWAVE into Madagascar.

## REFERENCES

- Brown, D., 1984, A note on the numerical solution of the wave equation with piecewise smooth coefficients: *Mathematics of Computation*, **42**, 369–391.
- Cohen, G. C., 2002, *Higher order numerical methods for transient wave equations*: Springer.
- Fehler, M., and J. Keliher, 2011, SEAM Phase 1: Challenges of subsalt imaging in Tertiary basins, with emphasis on deepwater Gulf of Mexico: *Society of Exploration Geophysicists*.
- Fomel, S., 2009, Madagascar web portal: <http://www.reproducibility.org>, accessed 5 April 2009.
- Hu, W., A. Abubakar, and T. Habashy, 2007, Application of the nearly perfectly matched layer in acoustic wave modeling: *Geophysics*, **72**, SM169–SM176.
- Levander, A., 1988, Fourth-order finite-difference P-SV seismograms: *Geophysics*, **53**, 1425–1436.
- Moczo, P., J. O. A. Robertsson, and L. Eisner, 2006, The finite-difference time-domain method for modeling of seismic wave propagation: *Advances in Geophysics*, **48**, 421–516.
- Muir, F., J. Dellinger, J. Etgen, and D. Nichols, 1992, Modeling elastic fields across irregular boundaries: *Geophysics*, **57**, 1189–1196.
- Padula, A. D., W. W. Symes, and S. D. Scott, 2009, A software framework for the abstract expression of coordinate-free linear algebra and optimization algorithms: *ACM Transactions on Mathematical Software*, **36**, 8:1–8:36.
- Santosa, F., and W. W. Symes, 2000, Multipole representation of small acoustic sources: *Chinese Journal of Mechanics*, **16**, 15–21.
- Symes, W., 2014, IWAVE structure and basic use cases, *in* TRIP 2014 Annual Report: The Rice Inversion Project. (available summer 2015).
- Symes, W. W., D. Sun, and M. Enriquez, 2011, From modelling to inversion: designing a well-adapted simulator: *Geophysical Prospecting*, **59**, 814–833. (DOI:10.1111/j.1365-2478.2011.00977.x).

Symes, W. W., and T. Vdovina, 2009, Interface error analysis for numerical wave propagation: *Computational Geosciences*, **13**, 363–370.

Terentyev, I., and W. W. Symes, 2009, Subgrid modeling via mass lumping in constant density acoustics: Technical Report 09-06, Department of Computational and Applied Mathematics, Rice University, Houston, Texas, USA.

Virieux, J., 1984, SH-wave propagation in heterogeneous media: Velocity stress finite-difference method: *Geophysics*, **49**, 1933–1957.

——, 1986, P-SV wave propagation in heterogeneous media: Velocity stress finite-difference method: *Geophysics*, **51**, 889–901.

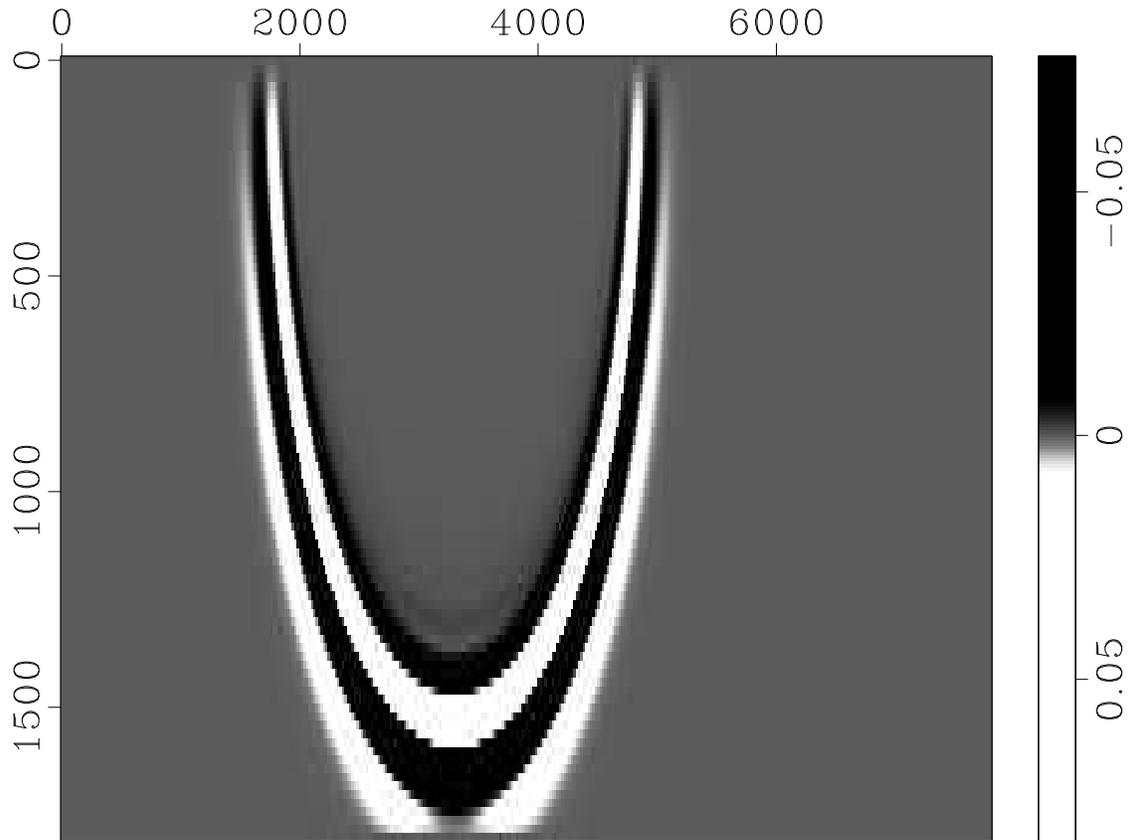


Figure 1: Point source field, homogeneous medium with  $v_p = 1.5$  km/s, at 1.2 s

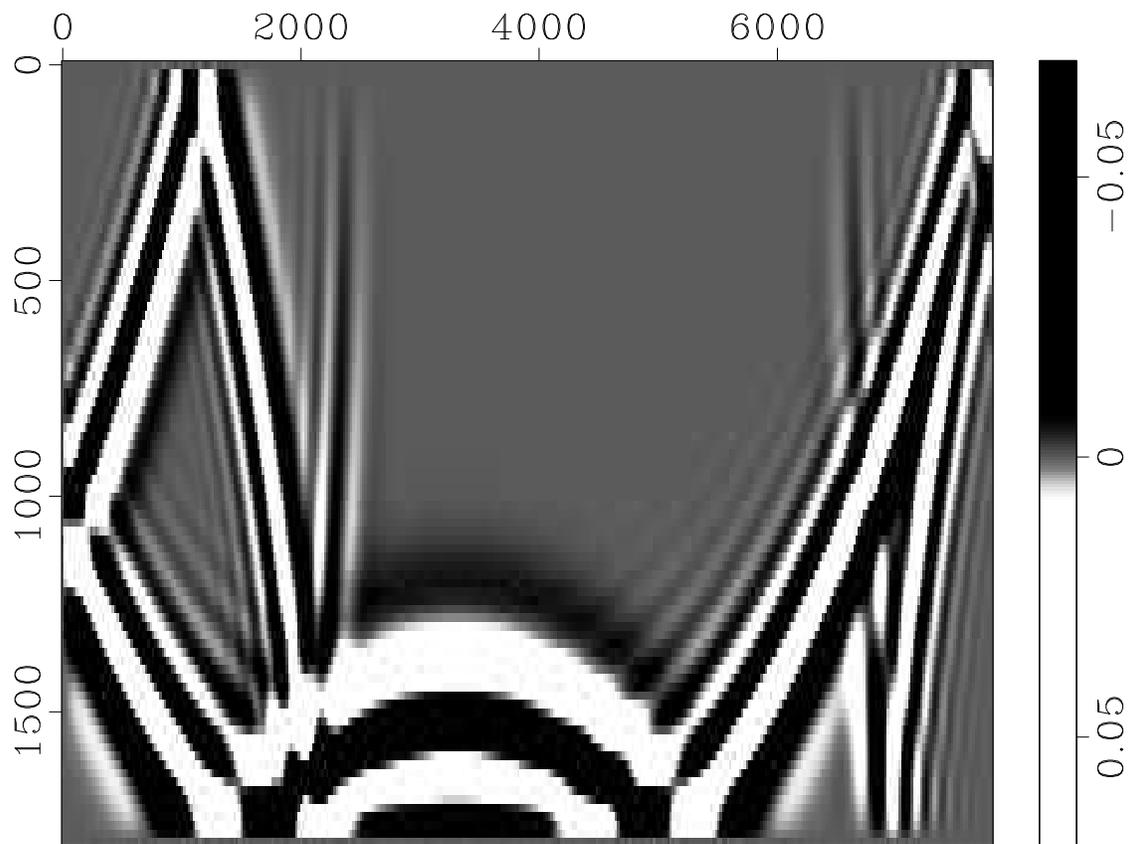


Figure 2: Point source field at 4.0 s, after interaction with reflecting (zero-pressure) boundaries

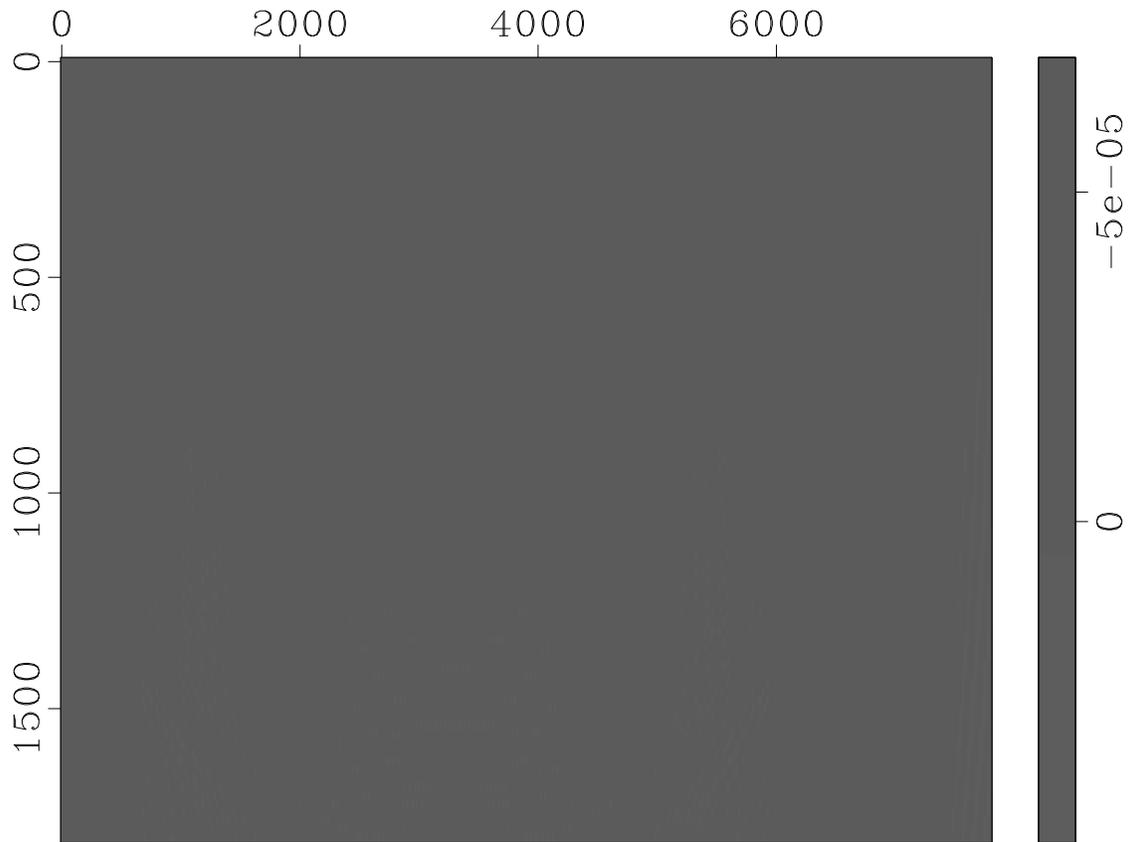


Figure 3: Point source field at 4.0 s, after interaction with 250 m PML boundary zones on bottom and sides ( $\eta_0 = 1.0$ ) - same grey scale as Figure 2. Longest wavelength carrying significant energy is roughly 500 m.

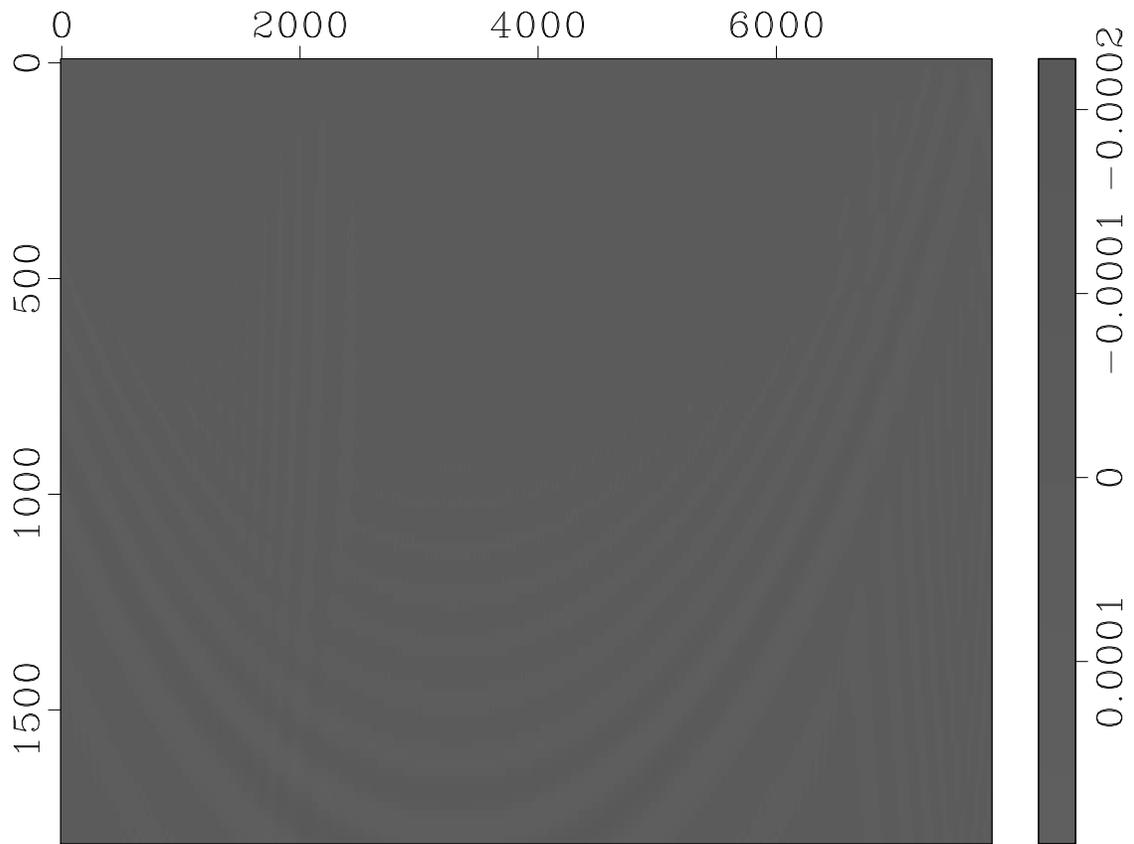


Figure 4: Point source field at 4.0 s, after interaction with 100 m PML boundary zones on bottom and sides ( $\eta_0 = 1.0$ ) - same grey scale as Figure 2.

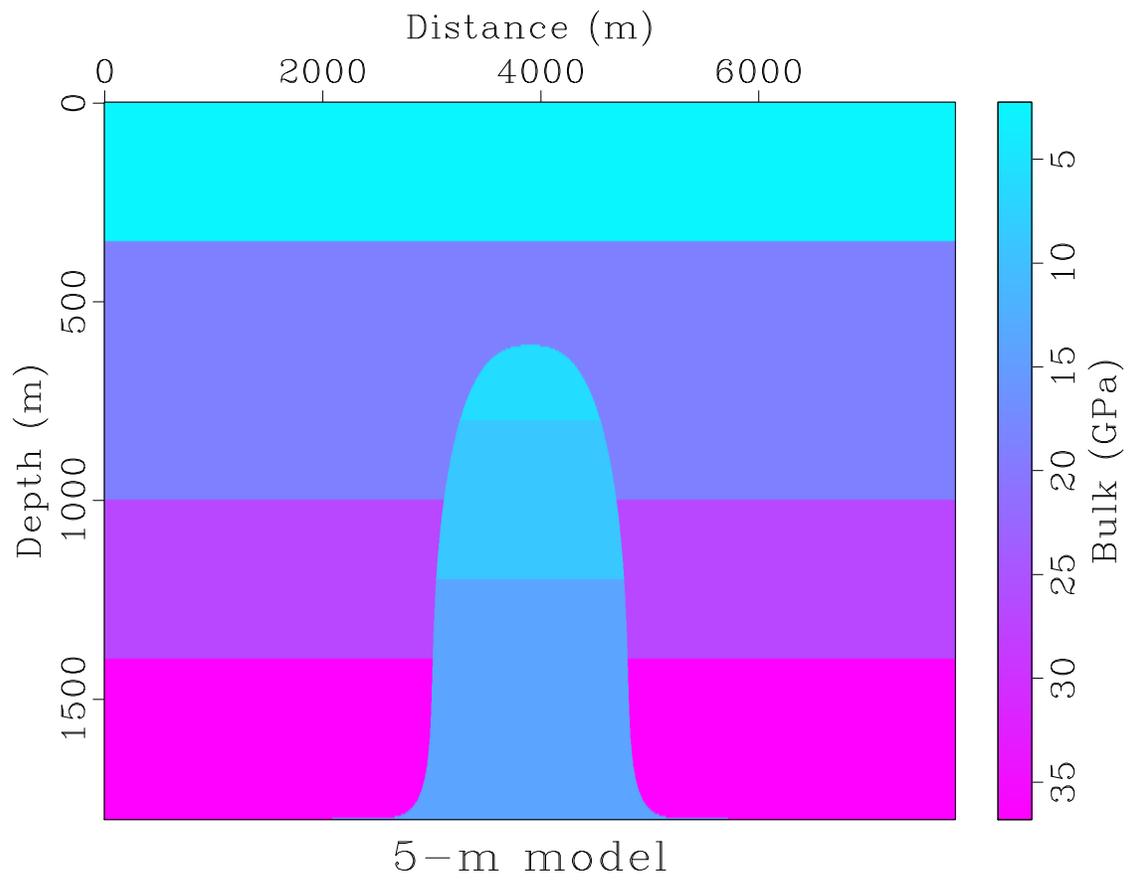


Figure 5: Dome bulk modulus

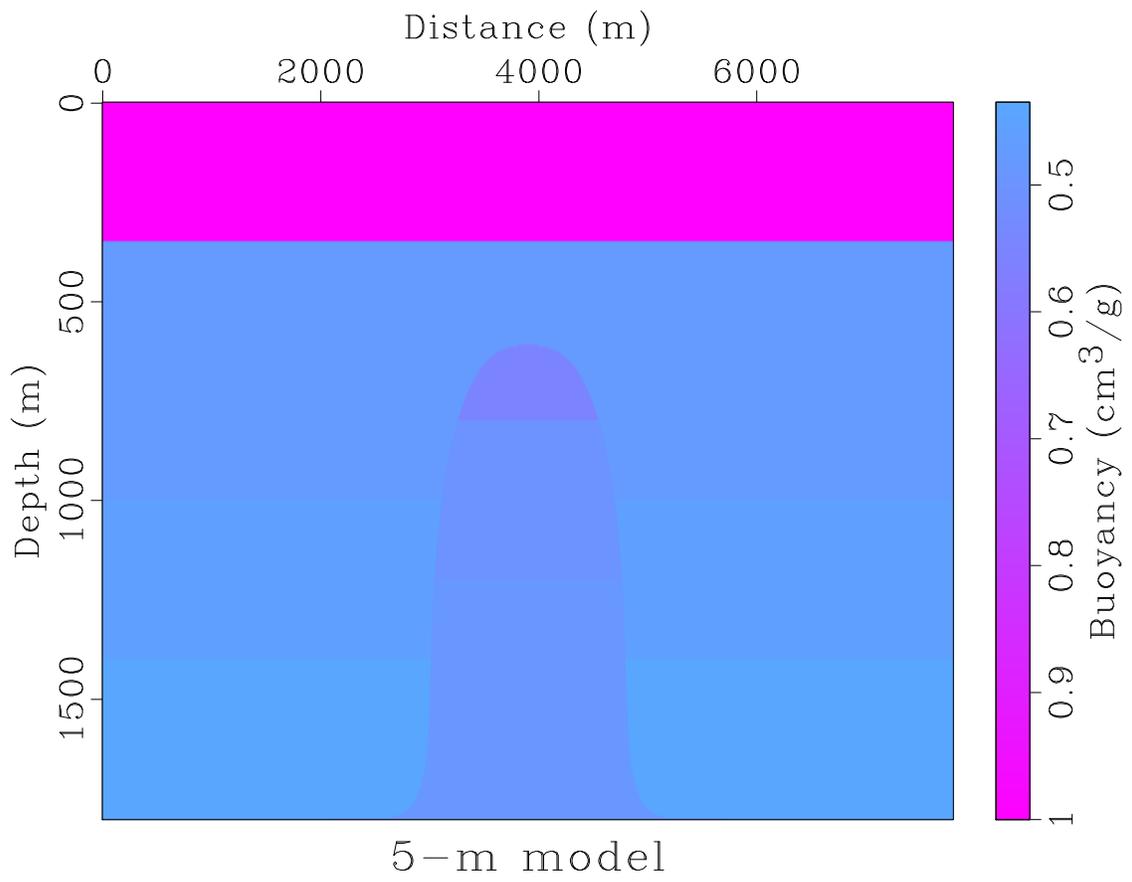


Figure 6: Dome buoyancy

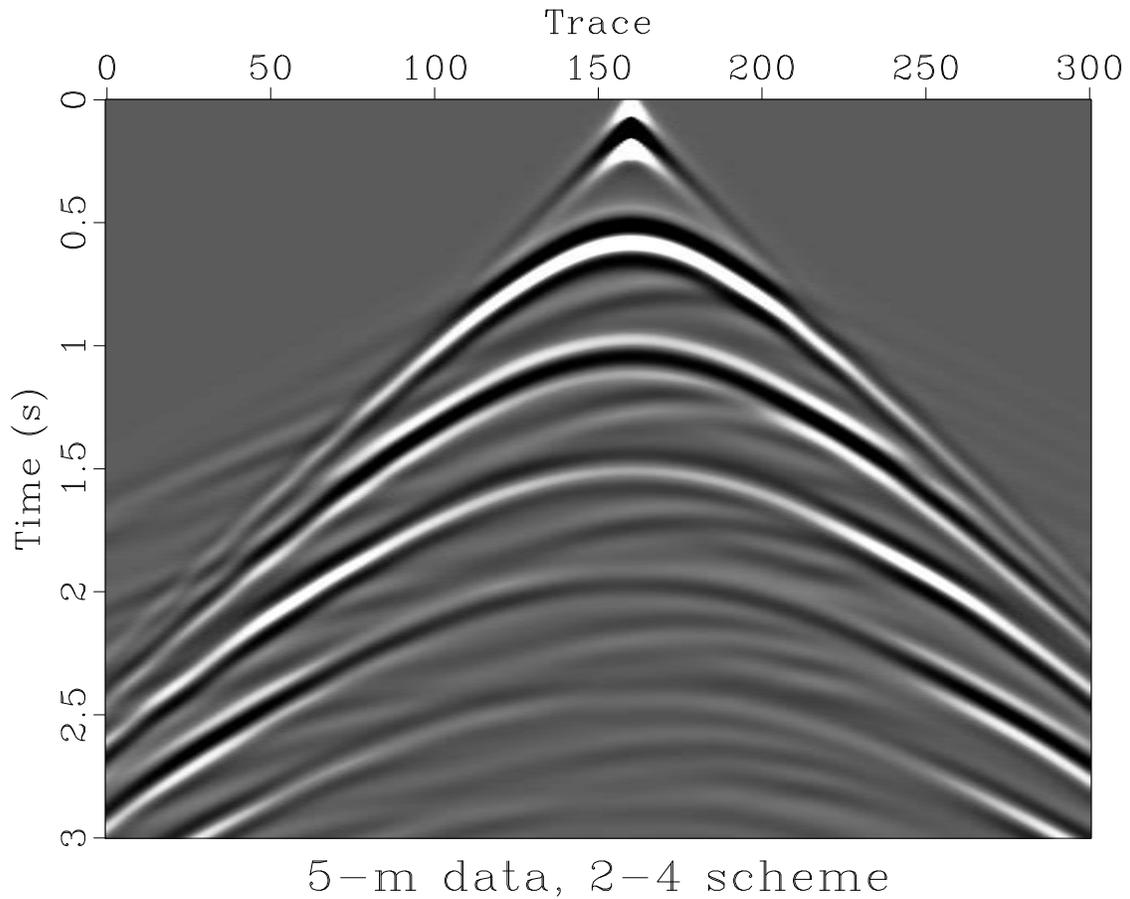


Figure 7: 2D shot record, (2,4) staggered grid scheme,  $\Delta x = \Delta z = 5$  m, appropriate  $\Delta t$ , 301 traces: shot  $x = 3300$  m, shot  $z = 40$  m, receiver  $x = 100 - 6100$  m, receiver  $z = 20$  m, number of time samples = 1501, time sample interval = 2 ms. Source pulse = zero phase trapezoidal [0.0, 2.4, 15.0, 20.0] Hz bandpass filter.

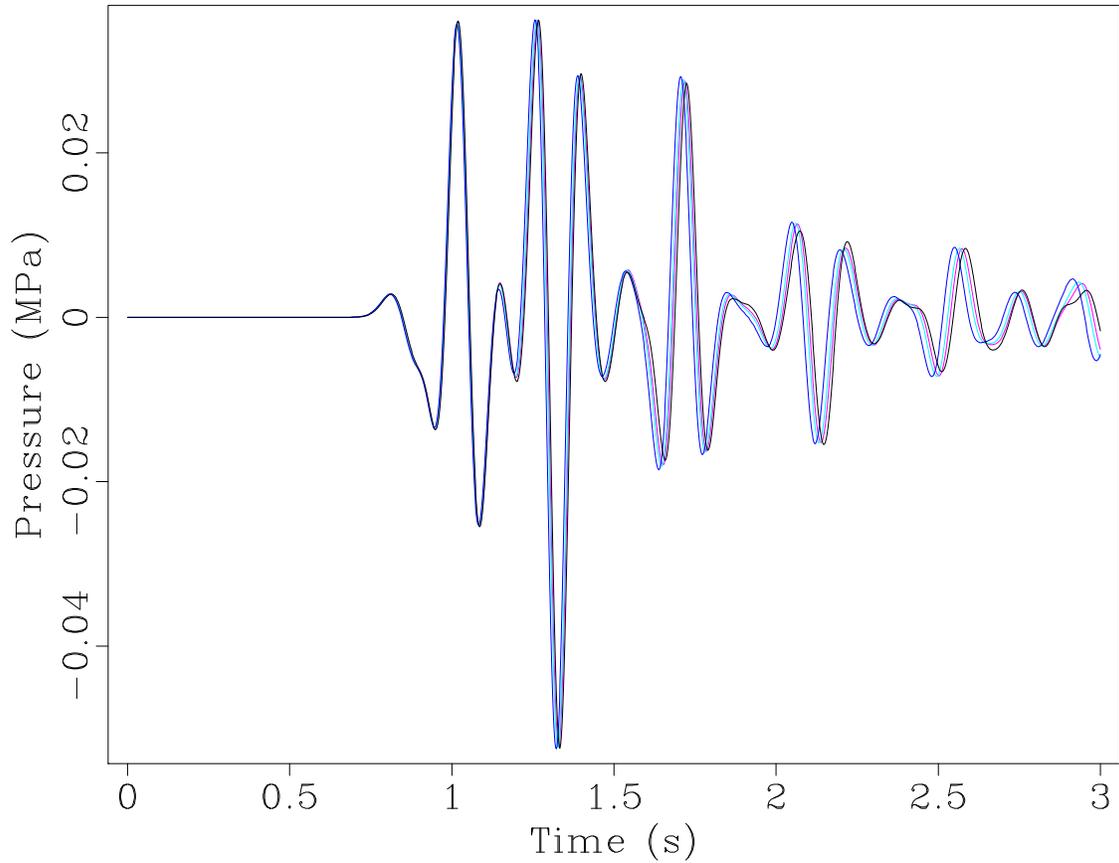


Figure 8: Trace 100 (receiver  $x = 2100$  m) for  $\Delta x = \Delta z = 20$  m (black), 10 m (blue), 5 m (green), and 2.5 m (red). Note arrival time discrepancy after 1 s: this is the interface error discussed in (Symes and Vdovina, 2009). Except for the 20 m result, grid dispersion error is minimal.

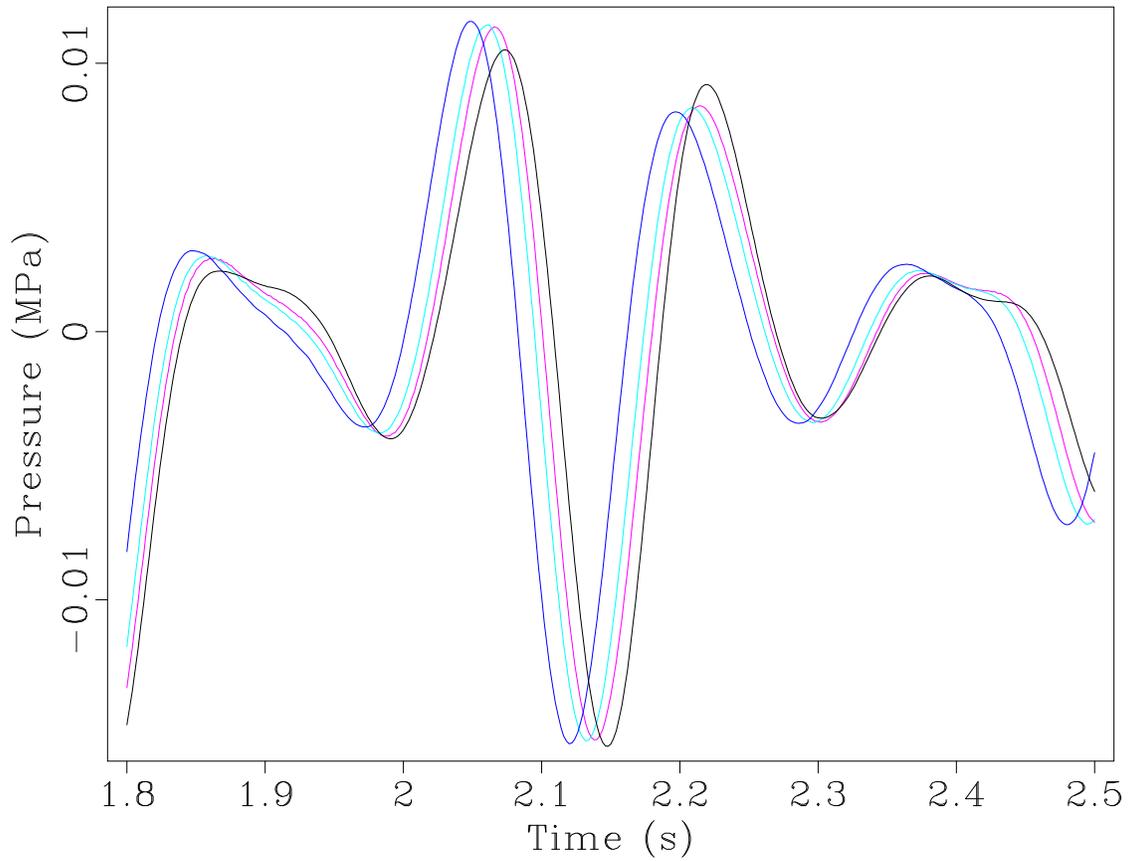


Figure 9: Trace 100 detail, 1.8-2.5 s, showing more clearly the first-order interface error: the time shift between computed events and the truth (the 2.5 m result, more or less) is proportional to  $\Delta t$ , or equivalently to  $\Delta z$ .

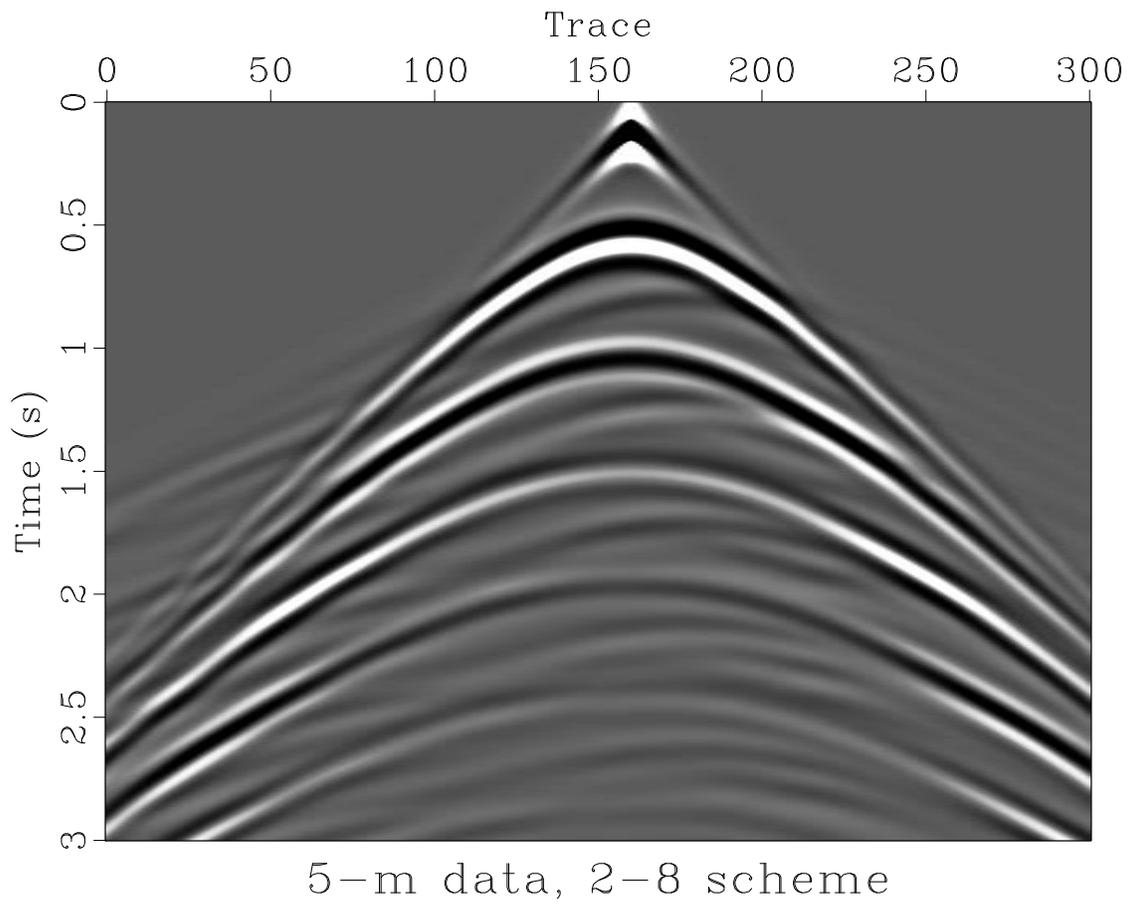


Figure 10: 2D shot record, (2,8) scheme, other parameters as in Figure 7.

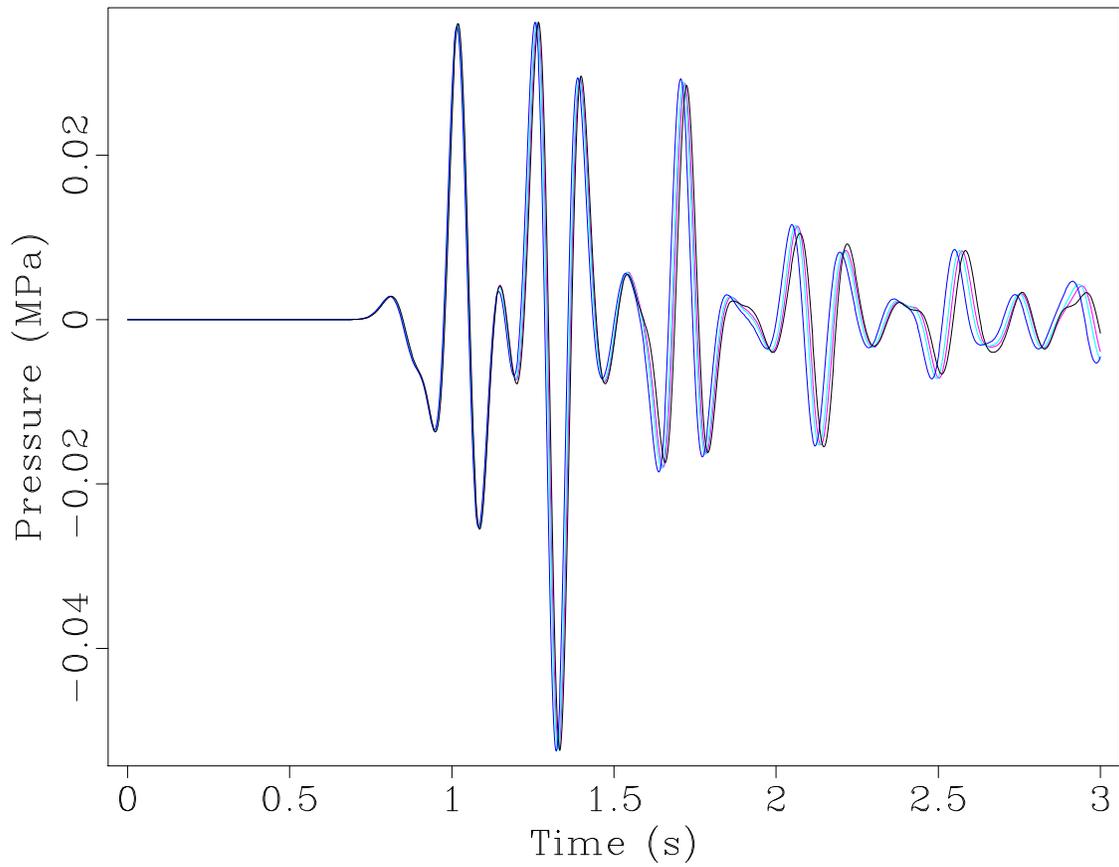


Figure 11: Trace 100 computed with the (2,8) scheme, other parameters as described in the captions of Figures 7 and 8.

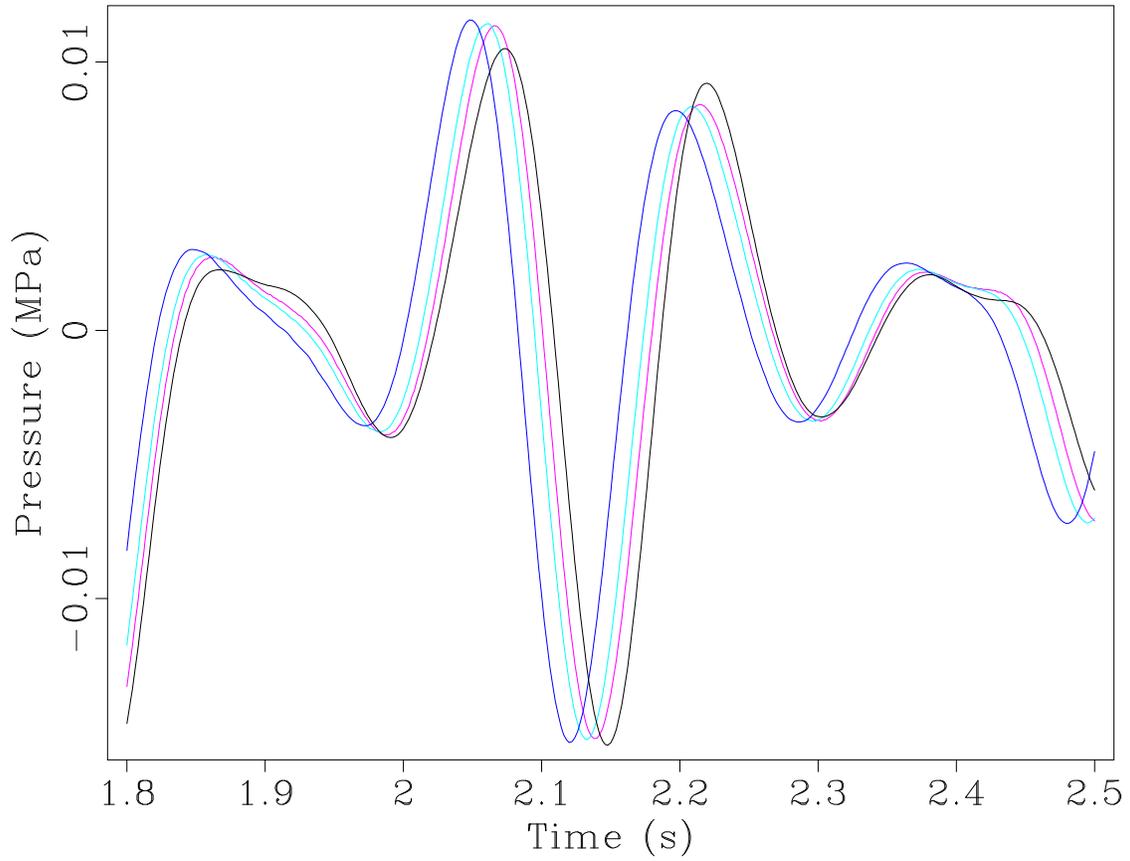


Figure 12: Trace 100 detail, 1.8-2.5 s, (2,8) scheme.. Comparing to Figure 9, notice that the dispersion error for the 20 m grid is considerably smaller, but the results for finer grids are nearly identical to those produced by the (2,4) grids - almost all of the remaining error is due to the presence of discontinuities in the model.