RI 9617

REPORT OF INVESTIGATIONS/1996

PB96-139712

3DTOM: Three-Dimensional Geophysical Tomography

UNITED STATES DEPARTMENT OF THE INTERIOR

BUREAU OF MINES

REPRODUCED BY: NTIS
U.S. Department of Commerce
National Technical Information Service
Springfield, Virginia 22161

U.S. Department of the Interior
Mission Statement

As the Nation's principal conservation agency, the Department of the Interior has responsibility for most of our nationally-owned public lands and natural resources. This includes fostering sound use of our land and water resources; protecting our fish, wildlife, and biological diversity; preserving the environmental and cultural values of our national parks and historical places; and providing for the enjoyment of life through outdoor recreation. The Department assesses our energy and mineral resources and works to ensure that their development is in the best interests of all our people by encouraging stewardship and citizen participation in their care. The Department also has a major responsibility for American Indian reservation communities and for people who live in island territories under U.S. administration.

Report of Investigations 9617

3DTOM: Three-Dimensional Geophysical Tomography

By Michael J. Jackson and Daryl R. Tweeton

UNITED STATES DEPARTMENT OF THE INTERIOR
Bruce Babbitt, Secretary

BUREAU OF MINES
Rhea Lydia Graham, Director

This report has been technically reviewed, but it has not been copy edited because of the closure of the agency.

CONTENTS

ILLUSTRATIONS

TABLE

UNIT OF MEASURE ABBREVIATIONS USED IN THIS REPORT

cm	centimeter	Mb	megabyte
dB	decibel	MHz	megahertz
ft	foot	m	meter
GHz	gigahertz	μs	microsecond
H	henry	mm	millimeter
Hz	hertz	ms	millisecond
in	inch	Np	neper
kB	kilobyte	s	second
kHz	kilohertz	S	siemens
km	kilometer		

DISCLAIMER OF LIABILITY

The Bureau of Mines expressly declares that there are no warranties expressed or implied which apply to the software contained herein. By acceptance and use of said software, which is conveyed to the user without consideration by the Bureau of Mines, the user hereof expressly waives any and all claims for damage and/or suits for or by reason of personal injury, or property damage, including special, consequential or other similar damages arising out of or in any way connected with the use of the software contained herein.

3DTOM: Three-Dimensional Geophysical Tomography

By Michael J. Jackson[1] and Daryl R. Tweeton[2]

ABSTRACT

3DTOM is a DOS-compatible computer program developed by the U.S. Bureau of Mines for three-dimensional tomographic imaging of the subsurface at mine sites. The program uses the simultaneous iterative reconstruction technique (SIRT) to invert travel-time data and produce maps of wave velocity, or to invert amplitude data and generate maps of wave attenuation coefficients. Either seismic (compressional and/or shear) or electromagnetic (e.g., radio or radar) wave data may be used. Ray tracing in 3DTOM uses several different methods, including ray bending, network theory, and a combination of these. User-defined constraints are important in reducing the mathematical nonuniqueness of inversions based on limited data. 3DTOM permits the use of hard constraints, or soft constraints based on fuzzy logic, to allow for uncertainty in the constraints. Reliable subsurface images are useful in many different mine-related problems, including void detection, fracture detection, fluid monitoring, and qualitative stress evaluation.

---

[1]Geophysicist, Department of Geology and Geophysics, University of Minnesota, Minneapolis, MN

[2]Research physicist, Twin Cities Research Center, U.S. Bureau of Mines, Minneapolis, MN

INTRODUCTION

This report describes the Bureau of Mines tomographic software 3DTOM, and consists of two principal parts: a user's guide and a technical overview. It is assumed that the reader is generally familiar with the fundamental concepts of tomographic imaging (summaries are presented by (*1*, *2*, *3*)[3]). The technical overview is presented for those who wish to understand the mathematical and physical foundations of geotomography. Although a good technical understanding is helpful for interpretation of results, the user's guide presents enough background for most users to begin using the program.

3DTOM differs from previous Bureau tomographic software (BOMTOM (*4*), BOMCRATR (*5*), and MIGRATOM (*3*)) primarily in its ability to perform inversions in three dimensions. However, it may also be used for two-dimensional imaging, and it incorporates several improvements in speed and stability for 2-D analysis. 3DTOM also provides more support for preliminary data analysis, and an integrated routine for contour plotting of the tomographic results. 3DTOM has a variety of methods for ray tracing, allowing the user to set relative priorities for execution speed and accuracy. Finally, 3DTOM has a new, more graphical user interface that supports both keyboard and mouse input. 3DTOM was written in Pascal, and compiled using Borland Pascal[4] version 7.0.

3DTOM USER'S MANUAL

SYSTEM REQUIREMENTS

3DTOM was designed to run on PC-compatible (DOS-based) microprocessors. A minimum of 640 kB of RAM and an 80386 processor are required, but more RAM (4 Mb or more) and a faster processor (80486 or above) are highly recommended. Although 3DTOM is a DOS application, it is highly recommended that it be run under a multitasking operating system such as Windows, Windows NT, or OS/2, because execution may take hours or even days to complete, depending on processor speed, size of the data set, and how 3DTOM is set up. A multitasking operating system allows you to continue using your computer for other tasks while 3DTOM is running. 3DTOM has been tested with MS-DOS versions 5.x and 6.x, with Windows version 3.1, Windows for Workgroups version 3.11, Windows NT version 3.1, and OS/2 Warp version 3. A full hard disk installation, including the executable program, graphics drivers, source code, and example files, requires approximately 2 MB of disk space.

Under DOS, the amount of system memory available for software and data is limited to 640 kB, even though the physical memory capacity of a computer may be much greater. 3DTOM was written and compiled with Borland Pascal version 7.0, which provides tools for accessing memory above the DOS 640 kB limit. The files RTM.EXE and DPMI16BI.OVL enable 3DTOM to store arrays in extended (XMS) memory, making it possible to work with

---

[3]Italic numbers in parentheses refer to items in the list of references at the end of this paper.

[4]Reference to particular commercial products is for purposes of information only and does not constitute endorsement by the USBM.

large numbers of ray paths and voxels. The amount of XMS memory available is detected automatically; if XMS memory is insufficient, 3DTOM will issue an error message and abort. To use XMS memory, you must run a driver (e.g., HIMEM.SYS) during boot-up (by including it in your CONFIG.SYS file). More information on XMS extended memory is available in the DOS manual.

Graphical displays in 3DTOM are generated through the Borland Graphical Interface (BGI). The necessary files are stored in the \3DTOM\BGI subdirectory. Supported graphics hardware includes CGA, EGA, VGA, Hercules, and IBM 8514. The type of graphics hardware present is detected by the software, and the appropriate driver is loaded. If you have a higher-resolution graphics card (e.g., SVGA), 3DTOM will use it at less than its maximum resolution, with one of the drivers listed above.

INSTALLING AND RUNNING 3DTOM

It is strongly recommended that you install and run 3DTOM using multitasking software such as Windows or OS/2. However, if prefer, you can install an run the software directly from DOS.

Installing and Running from DOS

Files on the distribution disk are compressed, and must be expanded by the installation program (INSTALL.EXE). Insert the distribution diskette into floppy drive A (or B if appropriate) and log onto that drive by typing A: <Enter> at the DOS prompt. Begin the installation program by typing INSTALL at the A: prompt. This will create a directory named \3DTOM on your hard disk and expand the files from the distribution disk to this directory. The files include executable and source codes (3DTOM.EXE and *.PAS), graphics drivers, example data files (DEMO1.3DD and DEMO2.3DD) and (if approprate) an update file containing information more recent than this publication (README.TXT). To run the program, log onto the 3DTOM directory using the DOS CD command: CD \3DTOM <Enter>. Then type 3DTOM <Enter> to begin execution. (You can omit the CD step if you add \3DTOM to your DOS search path - see the DOS manual for instructions).

Installing and Running from Windows

Insert the distribution diskette into floppy drive A (or B if appropriate). Select "Run" from the Program Manager "File" menu, and type A:SETUP <Enter>. The installation program will create a directory named \3DTOM on your hard disk and expand the files from the distribution disk to this directory. It will also set up a 3DTOM program group in the Windows Program Manager, containing program icons for 3DTOM, a 3DTOM configuration program, and a Windows help file. Any of these programs can be started by double-clicking its icon.

The program information file (PIF) 3DTOM.PIF contains information for Windows about how to run the program, including memory, display, and multitasking options. To modify any of these settings, run the PIF Editor. By default, Windows allows any DOS program to access up to 1 MB of system memory. Inversion of large data sets with 3DTOM will often require more than 1 MB; you can allocate additional XMS memory, up to the total available on your computer (3DTOM will usually not use all of the memory allocated to it). 3DTOM does not

use EMS memory, so the EMS should be set to zero. Display memory should be set to low graphics, and display usage to full-screen. It is very useful to enable background processing, which lets 3DTOM run as a background process while you execute other programs in the foreground. Under "Advanced" options, multitasking priority may be adjusted. If you find that your foreground programs are running too slowly with 3DTOM executing in the background, set 3DTOM's background priority to a lower number. Conversely, if your foreground programs execute very quickly, you can allocate more background processing time to 3DTOM, enabling it to run faster. Trial and error will quickly lead you to appropriate priority settings.

USER INTERFACE

The user interface for 3DTOM was designed and implemented using Borland Pascal's Turbo Vision object-oriented framework. Although it is a text-based DOS program, it is designed to have a "feel" similar to Windows programs or other programs with a graphical user interface. The program operates by responding to user-initiated "events", i.e., keystrokes or mouse clicks. A variety of control elements are provided for the user to select options or specify information, including dialog boxes with radio buttons and check boxes, and resizable text windows with scroll bars. These will be described in more detail in the relevant sections below.

3DTOM Display Screen

The screen display consists of three basic elements: a menu (along the top of the screen), a status bar (on the bottom), and a "desktop" area in between (Figure 1). The desktop is initially empty, and is used to display information in dialog boxes or text windows during program execution. The number at the upper right-hand corner of the screen shows the amount of available memory (both DOS and XMS); the number is updated whenever it changes, for example when memory is allocated to store data read from a file. Frequently used commands are tied to "shortcut keys" F1 through F10; some of these are listed in the status bar. Menu items and shortcut keys that are temporarily disabled are indicated in gray.

Text Windows

Many kinds of information are displayed in text windows, including input data, starting and final models, and on-line help. Each window has a title on top and a number in the upper right (Figure 2). These window displays have many useful properties.

Scrolling. The view can be repositioned vertically or horizontally by pressing the cursor control keys (arrows, PgUp, PgDn). Scroll bars along the right and bottom edges show the relative position of the current view. Mouse clicks on the scroll bars' terminal arrows have the same effect as the keyboard arrows; clicking the scroll bar above or below the position indicator has the same effect as PgUp or PgDn. In addition, the position indicator can be moved by a "click and drag" mouse operation.

Sizing and Positioning. Text windows can be resized and/or repositioned individually, and can also be cascaded or tiled together. Using a mouse, a window can be moved by a click-

and-drag operation on the window title; it can be resized by a click-and-drag on the "handle" in the lower right corner. From the keyboard, Control-F5 puts the current window in "move/resize mode": arrow keys will move the window, and Shift-arrow key combinations resize it by dispacing the lower right corner. Pressing Enter returns the window to its normal (display) mode. A window can be maximized (temporarily fill the whole desktop) by pressing the F5 key or by clicking the maximize control in the upper right corner. A maximized window can be restored to its "normal" size by pressing F5 or by clicking the restore control (upper right).

Selecting Windows. When multiple windows are open, only one at a time is active (i.e., scrollable), and some may not be visible (they may be obscured by other windows). There are a variety of ways to select a particular window. If it is visible, clicking a window with the mouse makes it active. Pressing the Alt key and a window number activates that window, making it visible if it was not. The F6 key can be used to cycle through all open windows, making each visible and active in turn.

Dialog Boxes

A dialog box is a particular type of text window that 3DTOM uses to get information from the user (Figure 3). This input may involve typing text or selecting options using other controls such as radio buttons and input lines (described below). Each control on a dialog box has a label, usually with one letter highlighted. A particular control can be selected by pressing the Alt key together with the highlighted letter key, or by clicking its label with the mouse. Repeatedly pressing the Tab key will cycle through all of the controls. Most dialog boxes include an OK button and a Cancel button. Clicking on OK or pressing Enter closes the dialog box and tells 3DTOM to accept the information; clicking Cancel or pressing Esc returns 3DTOM to the state it was in before the dialog was opened.

Visual Controls

Input Lines. These are subwindows on a dialog box where numerical or text items are displayed and may be edited by the user. When a text input is first selected, its contents are highlighted, indicating overwrite mode: the contents will be deleted and replaced when you begin typing the new information. If you wish instead to edit the default value supplied, begin by pressing an arrow key to enter edit mode: the highlighting will disappear, and the Del, Backspace and arrow keys can be used to edit the contents.

History Boxes. These are special input lines that remember values they have read previously, and they are marked by an arrow at the right end. When a history box is selected, pressing the down arrow (or clicking the arrow control) expands the input line to display a list of previously entered values (if any). One of these may be selected by a mouse click or by using the arrow keys.

Radio Buttons. These are groups of items, only one of which may be selected. The selected item of the group is indicated by a black dot in parentheses ( ) in front of it. When a different item is selected by a mouse click or Alt letter combination, the dot moves to the newly selected item.

Check Boxes. Options that can be turned on or off are indicated by check boxes. A check box appears as square brackets [ ]. An X in the brackets indicates that the option has been enabled.

File Dialog Boxes

A particular type of dialog box that is frequently used is a file dialog box (Figure 4). These contain three different types of graphical control elements: a text-history box, a scrollable window, and two control buttons.

Under the label "Name" is a text-history box. A default file wildcard specification or mask appears in the box, and is initially highlighted to indicate that this box is the active control. The text in the box can be edited using the left and right arrow keys, the Home, End, Insert, Delete and Backspace keys, and alphabetical keys. Entering a file name and pressing <Enter> or clicking "OK" will cause 3DTOM to read from or write to the indicated file. Entering a different file specification will cause a list of matching files to be displayed in the window below.

This control is termed a text-history box because it also retains a list of files selected during program execution. This list can be viewed by pressing the down arrow or by clicking on the downward-pointing arrow next to the box (Figure 4). Files in the list (if any) can be selected via the arrow keys or the mouse.

The second control element of the file dialog box is the window under the label "Files", in which is displayed a list of files in the current directory matching the file specification, with the first one highlighted. (The file information pane at the bottom of the dialog box shows the size and date of the highlighted file.) The file list may be "navigated" by use of the arrow keys or the mouse, and a file can be selected either by pressing <Enter> or by double-clicking.

Menus

The main program menu appears at the top of the screen (Figure 1). It is used to activate a number of "pull-down" submenus, each containing lists of available actions. Menu items may be selected by positioning the mouse cursor over them and clicking the left mouse button, or by holding down the Alt key and pressing one of the letters from the list above. Once made, a menu selection may be cancelled by pressing the ESC key. Some menu items are unavailable under certain conditions. An unavailable item is "grayed out," (the exact appearance may vary depending on the sort of display your computer uses.)

MENU STRUCTURE

The main menu can be used to activate the following"pull-down" submenus:
F: the file menu;
D: the data menu;
M: the model menu;
S: the setup menu;
R: run (begin the tomographic calculations);
W: the window menu;
O: the options menu;

H: help.

The following section provides a very brief description of what each menu item does. More complete descriptions can be found in the appropriate sections of this manual.

File Menu

Change Directory: Opens a dialog box enabling you to specify a directory for 3DTOM to use in reading/writing files. The initial default is \3DTOM.
Open Data: Reads a travel-time or amplitude data file.
Open Model: Reads a velocity or attenuation model.
Open Borehole Log: Reads a file containing measured velocity or attenuation values in a borehole, and modifies the current model to match the measured values.
Save Model As: Saves the current model.
Export Model: Saves the current model in a format that can be used by commercial graphics software.
Save Residuals: Saves the convergence history (RMS residual for each iteration) and a sorted list of largest residuals (difference between measured and calculated data).
Save Calculated T: Saves calculated travel times or reduced amplitudes.
DOS Shell: Temporarily exits to DOS.
Exit: Terminates 3DTOM.

Data Menu

Open File: Read a travel-time or amplitude data file.
Import: Read a MIGRATOM data file.
Save As: Save the current model in standard format.
Clean: Delete outliers from the current data set.
Close: Release the current data from memory.
View Info: Show a summary of the current data set (only available when display mode is set to text window.)
Display Mode: Display the current data set as either a text window or summary dialog box.
Analyze: Display time-distance plots, velocity histograms, and ray-path stereoplots for the current data set.

Model Menu

Open File: Read a velocity or attenuation model.
Import: Read a MIGRATOM model file (*.SUM).
Generate: Create a new grid and model.
Save As: Save the current model in standard format (Same as File Menu: Save As).
Export: Save the current model in a format that can be read by commercial graphics software (Same as File Menu: Export).
Extract: Examine the values at specified locations in the current model.
Close: Release the current model from memory.
View Info: Show a summary of the current model (only available when display mode is set to text window.)
View Grid: Show a graphical projection of the model grid.
Display Mode: Display the current model as either a text window or summary dialog box.

Contour Plot: Display a two-dimensional slice through the current model, with contoured velocity or attenuation coefficients.
Edit: Inspect or modify model values and node constraints.
Regrid: Change the node spacing in the grid, and interpolate values from the current model to the new grid.
Compare: Compute differences, ratios, or correlation coefficients for two models, or the average of several models.
Filter: Apply simple smoothing, horizontal smoothing, or spatial coherence filter operators to the current model.
Sampling Density: Generate statistics on raypath sampling density, or show a contoured display of sampling density.

Options Menu

Inversion Parameters: Set parameters for a tomographic inversion, including termination criteria, ray-tracing method, etc.
Coordinate System: Specify or transform coordinate system.
Units: Specify or convert units of measure.
Save As: Save current options in a file.
Retrieve: Load previously-saved options from a file..

Windows Menu

Resize/Move: Change the size or position of the current text window.
Zoom: Maximize the current window.
Next: Move the focus to the next window.
Tile: Make all windows approximately equal size and spread evenly over the screen.
Cascade: Overlap windows with the title visible for each.

Setup Menu

Mouse: Configure the mouse buttons and click speed.
Colors: Modify the default colors for visual controls.
Set BGI Path: Specify the location of graphics driver files.
Save: Save current setup to disk.
Retrieve: Retrieve setup from disk.

Help Menu

About: Show version number and release date for 3DTOM.
Notes: Display short help notes on specific topics.
View Help File: Open a scrolling window for the DOS help file.
Calculator: Display a simple 4-function calculator.

DATA

The data required by 3DTOM to carry out an inversion consist of source and receiver locations and measurements of travel time or reduced amplitude. The data must be supplied in a file with the proper format (see the section entitled Files).

Units

Units are arbitrary but must be self-consistent, and the units for the data determine those of the reconstruction. For example, with locations of sources and receivers specified in meters and travel times in milliseconds, the calculated velocities will have units of m/ms (or equivalently, km/s). With reduced amplitude units of nepers (Np) and spatial units of meters, the output has units of meters/neper, representing the attenuation distance (the distance over which the amplitude of a plane wave diminishes by 1/e). For more information on the latter see Amplitude Inversion in the Technical Overview. Units may be specified and converted through the Options menu.

Types

Seismic and EM data are treated interchangably by 3DTOM. The only difference is that the most appropriate units should be used for each (e.g., ms for seismic and μs for EM). In either case, amplitude data require preprocessing.

Analyzing Data

3DTOM contains three tools for basic preliminary data analysis: time-distance plots, velocity histograms, and stereographic ray-path projections. These are each based on the straight-line separation of source and receiver positions. These help to characterize the nature of the rock mass, and to identify outliers due to incorrect time picks, location errors or other problems with the data, prior to tomographic inversion.

3-D MODELS

Velocity and attenuation are modeled as continuous functions on a 3-D grid. 3DTOM requires a data set and an initial model in order to compute a final model of velocity or attenuation.

A default starting model is constructed by 3DTOM when a data file is opened. The default model can be replaced either by reading an appropriate model file or by generating a new model through the Model menu. After a model has been constructed or read by 3DTOM, you can view the grid geometry and evaluate the number and spacing of nodes. The grid viewer displays an orthographic projection of all nodes, as well as transmitter and receiver locations. The view can be rotated and rescaled as needed. 3DTOM's model generating routine allows you to specify the size and number of grid cells, and to specify uniform, layered, or checkerboard starting models. The Model Editor can be used to modify the properties of these models and to impose various constraints. A Borehole Log file containing measured velocity values in a borehole or elsewhere can be read by 3DTOM and used to modify an existing model. The regridding tool enables you to change the cell size in the grid

and interpolate values from the current model.

After carrying out the calculations specified in the Options menu, 3DTOM allows you to examine the resulting model or use it in various ways. You can save the final model to a file for later use as a starting model (allowing you to resume processing at a later time) . You can view contour plots of two-dimensional slices of the 3-D model. The Export menu option allows you to save the model in a format that can be used by commercial graphics software. The Extract menu item allows you to determine the reconstructed properties at a list of specified points (not necessarily grid nodes), or to view an arbitrarily-oriented slice. The Compare menu provides tools for computing differences, ratios, and correlation between pairs of models, as well as for averaging groups of models.

The 3-D Grid

The velocity or attenuation model is defined on a regular rectilinear grid (Figure 5). Values (velocity or attenuation distance) are specified at the nodes, and computed in the intervening cells (or voxels) by multiple linear interpolation. The grid is oriented with its edges parallel to the coordinate axes defined by the data.

Default Grid

On reading a data file, 3DTOM constructs a default grid based on the number of data points and the range of transmitter and receiver locations. The minimum and maximum values of X, Y and Z coordinates in the data file define the limits of the survey volume. In order to accomodate raypath curvature outside the survey volume, the default grid has dimensions exceeding by 20% those of the survey volume. The number of cells is linked to the number of data points by the following relation: $NX = NY = NZ = 2*Ndata^{0.3}$. The total number of cells ($N = NX * NY * NZ$) thus exceeds the number of data points, and the tomographic inversion is therefore underdetermined and results are mathematically nonunique. It is important to note that reducing the number of cells to less than the number of measurements does not guarantee uniqueness.

Examining the Grid

From the Model menu, select View Grid. 3DTOM will display an orthographic projection of the grid nodes and, if a data file is open, the transmitter and receiver locations. The initial view is parallel to the (negative) Z axis. An orientation diagram on the right shows the projected orientations of the positive X, Y, and Z axes. Above the orientation diagram are the direction cosines of the view angle (cosines of the angles between the line of sight and each of the coordinate axes; e.g., 1.0, 0, 0 indicates that you are looking in the direction of the positive X axis).

You can change the view angle by pressing the arrow keys: the up and down arrows rotate the view around a horizontal axis that is perpendicular to the line of sight, and the left and right arrows rotate around the vertical. The rotation increment is displayed at the top right of the screen. It is initially set to 15 degrees, but may be changed by pressing the space bar. Selectable increments are 5, 15, and 45 degrees.

You can rescale the drawing with the + and - keys (to increase and decrease magnification, respectively).

Modifying the Grid

If you wish to alter the grid geometry, use the Model Generator to create a new one,

or regrid to change the number/size of the cells while retaining current model values. This is typically necessary for two-dimensional data sets, for which the default relation NX = NY = NZ may be inappropriate.

Regridding

Regridding enables you to change the number and spacing of nodes,while retaining the values from the current model. The grid volume cannot be changed; in other words you cannot change the coordinates of the grid corners, but can only change the number of cells into which the volume is divided. New node positions are calculated and corresponding values interpolated, and the resulting model is written to a file; the model currently in memory is not altered.

Regridding allows you to evaluate rapidly the large-scale velocity structure with a coarse grid, and then resume processing on a finer grid to determine the finer-scale structure. Regridding also enables you to construct best-fit one-dimensional models (e.g., NX = NY = 2, NZ > 2 for a horizontally layered model) for use as starting models for 3-D inversion.

Note that the interpolation involved in regridding will generally result in some degree of smoothing (unless the number of cells in each direction in the new grid is an interger multiple of the number in the original grid, so that the new grid has nodes at the same locations as those in the original grid.)

If you wish to change the location of the grid (e.g., change the edge positions), you must generate a new model. Generating a new model does not preserve current model values.

Creating Models

The model generator allows you to create a new grid, and to set the model values in several ways. A dialog box lets you select from uniform, layered, or checkerboard models (with respectively, constant velocity or attenuation, values alternating vertically, and values alternating in all three directions.) A grid dialog box enables you to specify the range of X-, Y- and Z-values, and the number of nodes in each direction. A model dialog box allows you to specify the value (velocity or inverse attenuation) for a uniform model, or two values for a layered or checkerboard model. Layered and checkerboard models are very useful as starting points for forward modeling and resolution tests.

Keep in mind when generating models that the memory required is approximately 120 bytes per node. Thus a model with 22 nodes in each direction (10,648 total nodes) requires approximately 1.3 MB of memory. Additional memory is required for data and for the path-length matrix used in the tomographic calculations. When running in DOS, 3DTOM is able to use all available extended (XMS) memory. Under Windows, you may have to change the setup to make more memory available, by clicking on the Setup icon in the 3DTOM program group (see Windows documentation for further details.)

Modifying Models

3DTOM provides a model editor for viewing and modifying an initial velocity model, and for applying node constraints. A mathematically unique velocity reconstruction is independent of the starting model chosen; thus, modifying the initial model allows the user to test for nonuniqueness. Furthermore, when the solution is nonunique due to insufficient angular coverage of ray paths, it is important to incorporate other available site information

into the inversion. One way to do this is to apply constraints, as discussed in a subsequent section. An additional method is to incorporate the site information into the starting model; in general, the reconstruction will remain as close as possible to the initial model while still satisfying the travel-time data as closely as possible.

The model editor has three components. The node editor operates on one node at a time; you can navigate through the 3-D grid, examining and modifying the model value and constraint parameter at each node. The line editor operates on lines of nodes in the X, Y, or Z direction. Changes you make are applied to all nodes in a selected line. The plane editor operates on planes of nodes; each plane in the grid has a fixed value of X, Y, or Z.

Node Editor

The node editor dialog box displays the number and location of the current node, plus editable input lines containing the current model value (velocity or inverse attenuation) and constraint code. (The onstraints section gives details on constraint codes.) Navigation buttons on the right side of the dialog box provide the means to move around in the 3-D grid. There are two buttons for each coordinate axis (X, Y, Z), one for moving in the positive (Next) and one in the negative (Prev) direction.

The editor is initially located at the origin of the grid, node #1. To select a node for editing, click on the navigation buttons or press the Alt-key and highlighted navigation letters. Each time the new node number and location are shown.

When the desired node is reached, you can change its value (velocity or inverse attenuation) by clicking on the V label, by double-clicking on the value displayed, or by pressing Alt-V. The displayed value will be highlighted, indicating overtype mode, and you can enter the new value. The new values are stored as soon as you move to a different node, or if you Quit. The constraint code can be changed in the same way.

Line Editor

The line editor operates in a manner similar to the node editor. An orientation dialog box first asks you to specify the orientation of the line to edit. You can only edit lines parallel to the X, Y, or Z axes. Set the radio button for the orientation you want to select, and click OK or press the enter key.

The line editor dialog box then appears, with the line location specified by two coordinates, and navigation controls for changing those two coordinates. The third coordinate, not displayed, is parallel to the line of nodes to be edited. If all nodes in the line have the same value, that value is displayed in the value input line; otherwise the input line is blank. Similarly, the constraint code is shown only if it has the same value for all nodes in the line. The value and constraint code can be altered or set by clicking the input label (V for value, or Constraint) and typing in the desired quantity. The new value and/or constraint code are assigned to each node in the line as soon as you move to a different line or quit.

Plane Editor

Operation of the plane editor is similar to the other grid editors. An orientation dialog box first asks you to specify the orientation of the plane to edit. You can only edit planes perpendicular to the X, Y, or Z axes. Set the radio button for the orientation you want to select, and click OK or press the enter key.

The plane editor dialog box then appears, with the plane location specified by a single coordinate, and navigation controls for changing that coordinate. The second and third coordinates, not displayed, are parallel to the plane of nodes to be edited. If all nodes in the

plane have the same value, that value is displayed in the value input line, otherwise the input line is blank. Similarly, the constraint code is shown only if it has the same value for all nodes in the plane. The value and constraint code can be altered or set by clicking the input label (V for value, or Constraint) and typing in the desired quantity. The new value and/or constraint code are assigned to each node in the plane as soon as you move to a different plane or quit.

Comparing Models

An important part of interpreting tomographic results often involves comparing models. 3DTOM provides the following set of tools for model comparison. All of them require that the models be geometrically compatible, i.e., that they have exactly the same number and location of nodes. If you wish to compare incompatible models, you may be able to regrid one of them to make it possible to do so.

For each of the comparison tools, you must have a model currently stored in memory (either by reading from a file, generating, or computing).

Difference Tomograms

From the Model menu, select Compare and Difference. A dialog box will open; specify the file that contains the model that you want to compare with the current model. In the second file dialog box, give the name of the file where you want the results stored. This file will then contain a "difference model," with values at each node equal to the value of the current model minus that of the comparison model.

Difference tomograms can be useful for highlighting changes due to some causal effect, for example fluid injection. A data set collected before injection yields a reference tomogram, and a data set collected during or after injection yields a tomogram that reflects the same geology, but may be different due to the fluid. The difference tomogram will have much of the geologic "background" removed, and will highlight changes related to the injected fluid.

Ratio Tomograms

From the Model menu, select Compare and Ratio. A dialog box will open; specify the file that contains the model that you want to compare with the current model. In the second file dialog box, give the name of the file where you want the results stored. This file will then contain a "ratio model," with values at each node equal to the value of the current model divided by that of the comparison model.

Ratio tomograms can be useful for combining different types of information, for example P and S wave velocities. If you are able to identify both P and S arrivals in seismic traces, you can assemble a separate data set for each. Inverting the S arrival times will result in an image of S-wave velocity. Similarly, you can use the P data to generate a P-velocity model. A ratio model will then provide images of the ratio Vp/Vs, which can be useful for geologic interpretation (see Seismic Imaging for more details.)

Average Tomograms

Because practical limitations on raypath coverage often result in imperfect resolution, it is common for tomographic models to contain artifacts and remnants of the starting models used. It is sometimes desirable to invert a data set several times, using different starting models, and to average the results to obtain a more robust final model.

Select from the Model menu Compare and Average. A dialog box will open, showing a list of the models to be averaged; initially the list contains only the current model. To add a model to the list, click on the Add button; this will open a file dialog box from which you can select another compatible model file. You may repeat this as often as you wish; click on OK when you have finished assmbling the list of files. An output file dialog box will then open; enter the file name and path where you want the average model to be stored.

Correlation of Models

It is frequently of interest to quantify how similar two models are, for example in forward modeling studies. 3DTOM allows you to do this by calculating the linear correlation coefficient for the two ordered sets of values.

Select Compare and Correlation from the Model menu, and select the file that you want to compare with the current model. The variance for each model, the covariance, and the correlation coefficient will then be calculated and displayed in a dialog box.

Exporting Models

For presentation-quality graphics, 3DTOM will export models for rendering by commercial graphics software. Select Export from the Model menu, and choose the format from the dialog box provided.

In this release of 3DTOM the following export formats are provided:

Tecplot version 5
Tecplot version 6
Spyglass Slicer

The latter two programs have a variety of methods for representing 3-D models, including two-dimensional contour slices, 3-D contour surfaces, and rotation. Other commercial packages may be supported in future releases of 3DTOM, based on demand.

The Tecplot export files are text files that must be processed by Tecplot's PREPLOT utility. These files contain both spatial coordinates and model values. The Slicer export file contains only the model values in a text-matrix format that can be imported directly into the Slicer program.

Extracting Model Values

You can examine the current model value at any point or set of points within the model volume. 3DTOM provides three tools for doing this. The first provides the value for a single point, the second for a set of locations stored in a file, and the third generates a contour plot on an arbitrarily-oriented planar slice. From the Model menu, select Extract.

Value at a Point.

A dialog box with three input lines allows you specify a location anywhere within the grid volume. The inputs are initially set to the grid origin, and the model value and gradient are displayed. Gradient units are defined in terms of the data units, e.g., (m/ms)/m = 1/ms. To change a coordinate, press the Tab key or click on the label of the coordinate, and enter the new value. Click on OK, and the value and gradient displays will be updated for the selected point. Repeat for as many locations as you wish, then click on Quit to close the dialog box.

Values at a Set of Points from a File.

If you wish to obtain a list of model values at a set of points (for example, corresponding to a borehole), you must first create a plain text file containing the coordinates of the points. There should be one line in the file for each location, with X-, Y- and Z-coordinates delimited by spaces or tabs.

From the Model menu select Extract and Points From File. A file dialog box will open for you to specify the name and path of the file containing the coordinate locations. A second file dialog box will ask for the output file name and path. The same coordinate locations will be written to the output file, along with the model value (velocity or inverse attenuation) and gradient at each point.

Limitation: 5000 locations. If you have more than 5000 locations in the file, the remaining ones will be ignored.

Contour Slices.

If you wish to view a contoured slice that does not coincide with a plane of nodes in the grid, select Extract Slice from the Model menu. You will have to specify three points defining the planar slice to be viewed. In particular, the two axes for the contour plot will be defined in terms of the three points in the following way. Axis 1, displayed horizontally, will extend from point 1 to point 2; axis 2, displayed vertically, will extend from point 1 to point 3. The three points you specify thus determine the orientation and location of the planar slice through the model, and also define the area on that planar slice to be contoured.

Contour Plots

3DTOM provides two means of viewing contoured slices through the 3-D model. The first lets you view slices in pre-defined orientations and locations, with a minimum of effort. The second (described under Extracting Model Values) enables you to specify the orientation and location of any slice you wish to generate.

Press F8, or select Contour from the Model menu. A dialog box will allow you to choose from 3 slice orientations: perpendicular to the X axis (i.e., the Y-Z plane), perpendicular to Y, or to Z (Figure 6). You will be able to create contoured slices coinciding with planes of nodes in the desired orientation. Specify which of the slices in that orientation you want to see; slices are numbered from 1 to the number of nodes in the selected orientation. The default is through the center of the 3-D model.

A dialog box enables you to set the contour interval and range, and to control the placement of tic marks on the axes; the defaults are usually adequate for routine viewing. You can also select whether or not to use the true aspect ratio (height to width). The default is to use the true ratio (no vertical or horizontal exaggeration). If you turn this option off, 3DTOM scales the plot to fill the available space on your screen. A final option is to overlay the grid outline on the plot.

When the slice is generated (Figure 7), 3DTOM uses a right-handed coordinate system with the positive Z-axis downward, unless you specify otherwise (see Options). If things appear to be backwards and/or upside down, check this. If you have specified measurement units for your data (via the Options menu), 3DTOM displays derived units for the model (i.e. m/ms for data defined as meters and milliseconds.) By default, 3DTOM displays velocity or inverse attenuation. If you wish to display the reciprocal values (slowness or attenuation), change the display option in the Model menu.

You can move from one slice to adjacent slices (with the same orientation) by pressing

the + or - keys. The unique coordinate of the slice is displayed on the top of the screen. Press the Enter or Esc keys to close the graphic display.

Filtering Models

Two simple filtering operators are available: omnidirectional smoothing and horizontal smoothing. Each of these operators computes for each node an unweighted average of the nearest neighbors, including the node itself (9 for horizontal smoothing; 27 for omnidirectional). The smoothed model is written to a file, and the model in memory is not altered. The smoothing operators may be used repeatedly to obtain greater degrees of smoothing.

A certain amount of smoothing is inherent in 3DTOM's grid structure, where velocity or attenuation coefficients are represented as spatially continuous functions. Therefore smoothing is generally not necessary, although it may be useful in certain circumstances. Omnidirectional smoothing is most useful when you are working with fine grids and noisy data. Under these conditions, strong but spurious anomalies may arise in the reconstruction, and smoothing helps to diminish them. With coarser grids, however, smoothing may result in a loss of resolution. Horizontal smoothing is most appropriate for horizontally-layered rock masses.

Figure 8 illustrates the effects of these smoothing operators on a radar velocity tomogram. The rock mass consists of horizontally-stratified sandstones, limestones and shales, and the structure is evident in the unsmoothed tomogram (Figure 8a), on a 9 x 9 grid for the plane illustrated. Omnidirectional smoothing (Figure 8b) in this example nearly obliterates the layered structure by averaging together alternating layers of nodes with high and low velocities. Horizontal smoothing (Figure 8c) enhances the lateral continuity of the reconstruction.

Spatial Coherence Filter

Introduction

Noisy travel-time data typically result in "noisy" velocity tomograms, with localized strong velocity variations associated with travel-time outliers. Spatial coherence filtering (SCF) was designed by Zhou (*6*) to selectively remove strong, highly localized, spurious velocity anomalies, while leaving the more "coherent" (spatially continuous) velocity structure intact. Figure 9 shows an example.

The range of velocities in the reconstruction is divided into intervals. For each node, the spatial coherence is defined as the number of other nodes, in the same velocity interval, to which it is connected either directly or indirectly. A large group of contiguous nodes in a single velocity interval would thus have a high spatial coherence, whereas a single-node anomaly would have the minimum possible spatial coherence (one). A cutoff value of spatial coherence must be specified, and for each node with a spatial coherence less than the cutoff, the velocity is averaged with those of the immediately adjacent nodes.

Velocity Intervals

The idea of spatial incoherence hinges on the being able to characterize the velocities of certain nodes as "anomalous" or significantly different from those of neighboring nodes.

This is accomplished by specifying the number and bounds of the velocity intervals. For example, suppose the following represents velocity values for the nodes in a two-dimensional grid:

```
1.5  2.7  2.6  2.2  2.8  2.7  1.9  2.7
2.7  2.6  2.7  2.7  2.1  2.8  2.6  2.7
2.2  2.6  2.1  2.6  2.7  2.2  2.7  2.6
```

Specifying two intervals, separated by a value of 2.0, would cause two of the values to be placed in the first group ($V < 2.0$) and the rest in the second group. The group 1 nodes are isolated, whereas the group 2 nodes are all contiguous. The spatial coherence values would thus be:

```
 1   22   22   22   22   22    1   22
22   22   22   22   22   22   22   22
22   22   22   22   22   22   22   22
```

The two "outliers" are apparent. If instead the groups were defined by a velocity of 2.5, seven nodes would be assigned to group 1 and the remaining 17 to group 2. The spatial coherence values would be:

```
 1   17   17    3   17   17    1   17
17   17   17   17    3   17   17   17
 1   17    1   17   17    3   17   17
```

Here there are four single-node anomalies (spatial coherence = 1), and a contiguous low-velocity zone (3) separating two relatively uniform, higher-velocity areas (17). How the filter would treat these zones depends on the coherence cutoff value, as described below.

In 3DTOM, you may specify up to 5 velocity groups. Observe that for the example above, three velocity groups defined by velocities of 2.0 and 2.5 would result in the same spatial coherence values obtained with two groups defined by a velocity of 2.5. Before applying the SC filter, study contoured velocity plots to try to identify values that separate localized anomalies from neighboring areas. 3DTOM will display a histogram of node velocities before requesting velocity group definitions.

Note that specifying only one velocity group will cause the filter to become inactive: all nodes will be in the same group and therefore all will have the same coherence. Increasing the number of velocity groups will in general cause more nodes to have low coherence values, and will result in greater smoothing.

Cutoff Value

Nodes whose spatial coherence values are less than the specified cutoff are averaged together with their immediate neighbors. Setting the cutoff very low results in mininal filtering; for a cutoff of one, the filter is inactive. Setting a high cutoff value results in greater smoothing; when the cutoff is larger than the maximum coherence value, the filter is equivalent to simple smoothing. Selecting an appropriate intermediate cutoff value will selectively smooth the spatially incoherent noise. The best cutoff value must be found by trial and error, but Zhou (*6*) recommends starting with something like 20.

FILES

Input Files

Data Files contain travel times or reduced amplitudes, as well as source and receiver locations. These are the the only input files that are required.

Model Files contain the locations of grid nodes and the initial value of velocity or attenuation for each node.

Borehole Files contain values of velocity (or attenuation) measured in boreholes or elsewhere, together with spatial coordinates. The current model is modified to match the measured values, and nodes can be constrained to maintain the match.

Extract Location Files contain spatial coordinates for which you wish to obtain calculated velocities or attenuation coefficients.

Comparison Model Files are ordinary model files that may be compared with the current model.

Output Files

Model Files contain grid node locations and computed values of velocity and attenuation. They also contain additional information on inversion parameters, data, and raypath sampling density.

Data Files contain source and receiver locations, calculated and measured travel times, residuals, and apparent velocity. These are useful for forward modeling and for locating errors in the data.

Residual Files contain information on convergence (RMS residual as a function of iteration number) and a sorted list of the largest residuals.

Export Files contain the model information in a format readable by commercial graphics software for presentation-quality output.

Extracted Value Files contain the locations from the Extract Location file, along with the calculated value at each point.

Model Comparison Files contain the results of comparisons between models (differences, ratios, or averages).

Data Files

The data file must have the following format, which is similar to that used for MIGRATOM (*3*). Two header lines contain any information you wish to include for descriptive purposes, with a maximum line length of 120 characters, but no other format restrictions. These header lines are followed by the data, which are arranged in eight

columns (separated by one or more spaces or by tabs): a ray identifier (integer), the Cartesian (x, y, z) coordinates of the source and of the receiver for each ray, and the travel time of the first arrival. The units of distance and time in the input file determine the units of velocity calculated by 3DTOM; see the section on Options for more information on units. This format differs from that by MIGRATOM in only one respect, namely the restriction that ray identifiers here must be numeric, rather than alphanumeric strings. This restriction removes the requirement for a fixed field width for the identifier, and generally makes it much easier to produce correctly readable data files. If you have existing MIGRATOM data files, you can read them into 3DTOM by selecting Import from the Data menu.

To keep your files and projects organized, and for easy accessibility, it is recommended that you set up a subdirectory under \3DTOM for each data set or group of related data sets, and all associated model files and other files.

Model Files

Model files must have the following format. An arbitrary number of lines containing descriptive text is optional. A line containing the character @ must be included after any header lines (even if there are none) and before the grid description lines. The four grid description lines contain:

Xmin Xmax
Ymin Ymax
Zmin Zmax
NX NY NZ

where Xmin and Xmax indicate the minimum and maximum X coordinates (and similarly for Y and Z), and NX, NY, and NZ indicate the number of nodes in the respective coordinate directions.

After the four grid description lines, there must be a line for each node in the grid, with X, Y, and Z coordinates and model value (velocity or inverse attenuation), delimited by spaces or tabs. Coordinates and model values must be in the same system of units as the data. The lines in the file should be arranged with X coordinates cycling in the innermost loop and Z coordinates in the outermost, e.g.:

```
@
1 3
1 3
1 5
3 3 5
        1       1       1       4.5
        2       1       1       4.4
        3       1       1       4.6
        1       2       1       4.4
        2       2       1       4.5
        3       2       1       4.4
        1       3       1       4.7
        2       3       1       4.6
        3       3       1       4.4
        1       1       2       4.3
```

etc. The Z-axis is typically vertical, and the Z coordinate may increase with either depth or elevation (see Coordinate System Options).

Nothing further is required when a model file is read. However, when 3DTOM writes a model output file, it includes the following additional information: the constraint code for each node; the raypath sampling density (number of rays that intersect each cell); and a list of unique source and receiver locations.

Two-dimensional models written by MIGRATOM (*.SUM) can be imported via the Model menu. If the original 2-D grid has constant values for the X or Y coordinate (i.e., if the plane is perpendicular to the X or Y axis), 3DTOM will generate a pseudo-2D grid one voxel wide (two nodes wide), with the same number of cells as the original grid and twice the number of nodes. If the original grid is oriented obliquely to the X and Y axes, a fully 3-D grid will be constructed.

Residual Files

A residual file contains the RMS residual for each iteration, plus a sorted list of the 50 largest residuals for individual datum points. The default file extension is ".RSD".

A residual is defined as the difference between a measured datum (travel-time or amplitude drop) and the corresponding value calculated for the current model. The root-mean-square (RMS) residual is the square root of the mean of the squared residuals. It indicates how well the model corresponds to the measured data. RMS residuals typically decrease in a quasi-exponential manner as iterations proceed and the model is progressively modified to better match the measured data.

A residual summary file (*.RSD) contains two types of information. First, the convergence history of the iterative reconstruction is documented, with the RMS residual as a function of iteration number. This file also lists the RMS velocity perturbation (indicating how much the velocities have changed from their starting values) after each iteration. Second, a sorted list is presented of the fifty largest residuals. This list can help identify problems in the data such as typographical errors or incorrectly picked travel times.

There is no residual file written by default. If you wish to save one, use the File menu.

Borehole Files

Introduction

Because of the underdetermined or ill-conditioned nature of geophysical tomography, it is generally desirable to: 1) use the best possible starting model, incorporating as much site information as possible, and 2) constrain the inversion to match any known boundary values. 3DTOM helps you to do this by reading borehole log files, modifying a model accordingly, and allowing you to set node constraints forcing the final solution to match the measured borehole values.

Borehole File Format

A borehole file contains one line for each measured value of velocity or inverse attenuation. Each line should contain the X, Y, and Z coordinates where the value was

determined, and the value, delimited by spaces or tabs. Note that the locations need not be in a borehole; they may be in multiple boreholes and/or anywhere else that velocity or attenuation measurements have been made.

Applying Borehole Information

Rather than manually editing a model node-by-node to match the borehole data, you can let 3DTOM do it automatically. You decide how to modify the model to take the borehole data into account. There are three ways to do this:

Nearest Neighbor - Here the borehole values only affect the nearest grid nodes; nodes more than one cell distant from the nearest borehole measurement are unaffected. For each affected node, the velocity is simply set equal to that of the nearest borehole measurement.

Inverse Distance - Unlike the nearest neighbor option, each borehole measurement here affects the model values at all nodes in the grid. For the inverse distance option, each node is assigned a weighted average of all measured values (with weights being exponential functions of inverse distance).

Polynomial Fit - This modifies the velocity or attenuation value for every node in the grid, by fitting a 3-dimensional polynomial to the measured values (*7*), using the method of singular-value decomposition (*8*). You can specify polynomial degrees of 0 (producing a constant-value model), 1 (producing a constant-gradient model), or 2 (producing a constant-second-derivative model). NOTE: You should carefully inspect models involving polynomial degrees greater than zero, because trends in the measured values can be extrapolated to unreasonable values, especially when the measured value locations are limited.

The different results of these methods are illustrated in Figure 10, using hypothetical P-wave borehole velocity logs for the DEMO1 data geometry. In addition to editing the model values, 3DTOM will also set node constraints for you. After reading the borehole file, a dialog box lets you select the application method from the list above, and lets you specify a constraint code that is applied to the same nodes.

You can use the methods in combination to set both model values and node constraints. Use the inverse-distance method or fit a polynomial first (modifying the entire model) while leaving the constraint set to zero or a low value, and then re-read the log file and use the nearest-neighbor option to modify the neighboring nodes and set strong fixed-boundary-value constraints.

CONSTRAINTS

Uniqueness and the Need for Constraints

Due to imperfect angular coverage in common source-receiver geometries, travel-time data may not contain sufficient information to obtain a mathematically unique reconstruction of velocities (*9*, *10*, *3*). It is important in such cases to use all other

available information about the site. Application of constraints enables you to eliminate velocity solutions that are mathematically acceptable (in terms of satisfying the travel-time data) but geologically unreasonable. For example, a nonunique inversion may result in a model with unrealistically high or low velocities, one that agrees poorly with measured borehole velocities, or a model velocity structure that appears unreasonable based on regional stratigraphic or structural information.

In 3DTOM, a distinction is made between constraints that are applied over the entire grid (global constraints) and those that are applied locally to portions of the grid (node constraints). These will be discussed in turn in the following sections.

Global Constraints

The global constraints that 3DTOM allows the user to specify are upper and lower bounds on reconstructed velocities. These bounds may be applied based on a general knowledge of rock types present at the site and physical considerations. For example, compressional velocities in crustal rock rarely exceed 6.5 m/ms, and even in intensely fractured rock they are not less than the speed of sound in air (0.3 m/ms). Electromagnetic wave velocities in rock are always less than in air (300 m/μsec). Tighter bounds than these may be reasonable in many cases.

Node Constraints

Node constraints are useful for forcing the solution to match known boundary values (e.g., measured borehole velocities), and for maintaining uniformity over groups of nodes (e.g., continuous layers). Because the latter type of constraint is often quite uncertain, node constraints are assigned an estimate of uncertainty, which is used as a weighting factor so that better-determined constraints have a stronger influence on the solution than do poorly-determined ones. Such weighted constraints are herein referred to as "fuzzy" constraints.

Node constraints may be entered manually, by means of the grid editor, or automatically, by reading a borehole file. Each node in the grid has an associated constraint parameter, which is a real number consisting of an integer part and a fractional part. The integer part designates the type of constraint to apply: 0 means the velocity is unconstrained and can be freely adjusted by 3DTOM to obtain the best possible agreement with the travel-time data; a negative number indicates that the velocity is to be held fixed at the value in the starting model (this is useful for forcing the solution to match measured borehole velocities); a positive number assigned to a selected group of nodes indicates that uniformity of velocity is to be maintained over that group (this is accomplished by averaging velocities over the group, and is useful, for example, for producing uniform layers in the reconstruction). Different groups may be selected and assigned different positive constraint parameters; uniformity is then enforced for each group, independently of the other groups.

The fractional part of the constraint parameter represents the uncertainty in the constraint, or the "fuzz factor" f ($0 \leq f < 1$). After each iteration, two different velocity values are calculated for each node: an unconstrained velocity v_0 determined by the SIRT equations; and a constrained velocity v_1 calculated by strict enforcement of the constraint. The fuzzy constraints are applied by replacing the unconstrained corrected velocity at each node with a linear combination of the unconstrained and fully constrained values ($v = v_0 * f$

$\pm v_i*(1-f)$ for $f<>0$). Thus larger f values correspond to greater uncertainty or fuzziness in the constraint.

OPTIONS

Units of Measure

The units for spatial coordinates and lengths can be defined as mm, cm, m, in, or ft, and converted from one to another. Data units can be specified as μs, ms, or s for travel times or dB or Np for reduced amplitude. Specified units are used in graphic displays.

To convert length and/or data units, select Units and Convert from the Options menu. A dialog box will show the current units (if they have been defined) on the left side, and available choices on the right side. If you have not previously defined the units for your data set, you can specify those on the left side. When both sets of units are defined, select OK. A confirmation dialog box will verify your conversion choices. A new data file will be written, containing the rescaled data. The data in memory are not affected, so if you wish to use the rescaled data, you must open the new data file. NOTE: unit conversion applies to data only; models are not rescaled. Therefore if you are going to convert units, you should do so before calculating models - once a model is computed and saved, the units cannot be changed.

Coordinate Systems

3DTOM assumes a Cartesian right-handed coordinate system with the positive Z-axis downward. If a different coordinate system was used, some of the graphic displays will be inverted (upside-down and/or backwards), unless you specify the coordinate system used.

Like unit conversions, coordinate transformations are applied to data but not to models, and the transformed data are written to a new data file without overwriting the data in memory. Coordinate transformations involve translation (displacement) and/or rotation. Translation of the coordinate sytem simply involves adding an offset value to each of the coordinates in the data set. Specify the values in the translation dialog box. For rotation, you specify the orientation of the rotation axis as a Cartesian vector (e.g., 0,0,1 for the Z axis), and the rotation angle. The rotation axis is defined by the origin of the coordinate system and the specified axis orientation. Rotations are clockwise when viewed parallel to the rotation axis. For example, suppose we have a cross-borehole data set, with two vertical boreholes having respective horizontal coordinates of (0,0) and (10,10), in a right-handed coordinate system where the X axis corresponds to north, Y to east, and Z to down. Hole 2 is thus northeast of hole 1. Rotation of 45 degrees around the +Z axis (0,0,1) will produce a new coordinate system with the X axis pointing northeast, and the horizontal coordinates for hole 2 will be (14.14,0) in the new data set (hole 1 coordinates are unchanged in this example.)

In the preceding example, the coordinate system rotation is quite advantageous for computational efficiency, because the data set is two-dimensional. Since the grid is always constructed with its edges parallel to the coordinate axes, two-dimensional data sets can be covered with a pseudo-two-dimensional grid (one voxel wide) only when the data plane is normal to one of the coordinate axes. The original data set in the example, prior to rotation, would require a fully three-dimensional grid (more than one voxel wide in the X, Y, and Z directions), in which most of the cells would be unsampled by ray paths, and for which processing would be less efficient.

Inversion Parameters

A dialog box enables you to set the parameters that control how 3DTOM computes a tomographic image.

Global Constraints.

Set the maximum and minimum allowed velocity or attenuation distance. 3DTOM will not allow values in the reconstruction to exceed the specified values.

Termination Criteria.

Maximum allowed number of iterations and residual tolerance. Iterations end when the RMS residual decreases to the specified tolerance or at the set maximum number, whichever occurs first.

Ray Tracing.

Number of iterations performed using straight rays and crooked rays. Remaining iterations (until first termination criterion is met) use curved rays, either by ray-bending alone or by a hybrid technique. You can also set the retrace interval, the number of iterations to perform for each ray tracing. Since ray tracing takes most of the execution time, setting the retrace interval to a value greater than 1 can greatly speed up a run.

INVERSION

Straight and Curved-ray Processing

Straight Rays.

This is by far the fastest method, but also the least accurate when velocity contrasts are large. As a rule, it is recommended for a quick first look for new data sets, or whenever velocity contrasts are less than about 20%. This method should also be used in general for amplitude inversions, because ray-path curvature is controlled by the velocity distribution rather than by attenuation. However, for EM amplitude inversion in the low-Q regime, slowness is proportional to attenuation, so curved rays may be used in that case (see the Technical Overview section on EM inversion for more details).

Network Theory Minimum-Time Paths.

This method is the most time-consuming, and execution time increases rapidly with model size. For each source location, it constructs a minimum-time "tree" or network of straight-line path segments from node to node throughout the grid. Each ray path is thus approximated as a sequence of line segments of size comparable to the cell dimensions, and these are here referred to as *crooked rays*. For more information see Moser (*11*) and the Technical Overview.

Ray Bending.

The bending method iteratively modifies a trial path (by default, a straight line) from source to receiver. The path is progressively subdivided into smaller intervals, and the interval endpoints are displaced according to the local velocity gradient. The processes is continued until a stable minimum travel time is obtained. The resulting ray paths are typically much

smoother then those obtained by the network method, and they are here referred to as *curved rays*. For a complete description of the method, see Um and Thurber (*12*) and the Technical Overview.

Hybrid Method.

In some cases (for example, strongly stratified velocity variations) the bending method may converge to a local rather than to a global minimum travel time. In other words, it may find a ray path corresponding to an arrival later than the first arrival. This problem can be averted by using crooked rays (estimating the global minimum-time path) rather than straight lines as the starting point for ray bending. This approach requires the most processing time, but is the most accurate.

SIRT and RAYPT

The simultaneous iterative reconstruction technique (SIRT) is the basis for 3DTOM. The Ray-projection technique (RAYPT) was developed by Singh and Singh (*13*). It is a constrained optimization technique designed for imaging discrete anomalies in a uniform background material.

Before using RAYPT, you should examine the data carefully, especially time-distance plots and velocity histograms. These must serve as the basis for your interpretation. You have to estimate the property (velocity or attenuation) value for the background material. If the material is indeed uniform with isolated anomalies, the background material may correspond to a peak in the velocity histogram. You may also be able to observe a tight linear trend on the time-distance plot (corresponding to rays through the background) with groups of data points above and/or below (corresponding to rays that sampled anomalous regions.) You must decide what range of apparent velocities correspond to the background; rays with apparent velocities outside this range are then considered as having passed through one or more anomalies.

When you have finished this interpretation, open the Inversion Parameters dialog box and click on RAYPT; set the other parameters as needed, and then click on OK. A dialog box will open for you to specify the value characteristic of the uniform background, and a range. For example if your interpretation has rays with apparent velocities between 4.0 and 4.3 passing through the uniform background, and other rays sampling both background and anomalous material, then set the background value to 4.15 and the +/- range to 0.15.

Then run the inversion. Calculations are carried out as usual, but with additional constraints imposed by RAYPT, tending to make the model agree more closely with your interpretation. Figure 11 illustrates the difference between SIRT and RAYPT reconstructions using synthetic data computed for a target model with a discrete high-velocity anomaly in a uniform medium. The RAYPT reconstruction more accurately indicates the sharpness of the anomaly and the uniformity of the background.

A Step-by-step Guide to Tomographic Inversion with 3DTOM

Necessary and optional steps for producing a tomographic image for a travel-time or amplitude data set are outlined in this section.

Create data file.

This is most easily accomplished using any commercial spreadsheet software. The data format is described in the Files section. For travel-time data, no pre-processing is required. For amplitude data, you must carry out the data reduction steps outlined in the section on Files. Export or save the spreadsheet as a tab- or space-delimited ASCII text file, using a file name ending with '.3DD', the default data file extension for 3DTOM. Create a subdirectory of \3DTOM for this data set and associated files that will be produced by 3DTOM, and put the data file in this subdirectory.

Create borehole log file.

If you have measured values of velocity or attenuation in boreholes or elsewhere, and you want to include these values in the tomographic inversion, create a borehole log file as described under Files. Use a filename extension of '.LOG', and save the file in the same subdirectory as the data.

Run 3DTOM.

Because the tomographic calculations may take a substantial length of time to execute, it is recommended that you run 3DTOM in a multitasking environment such as Windows or OS/2. Select a full-screen display in order to be able to view 3DTOM's graphic displays.

Load the data file.

Press F3, or click on the F3 label on the bottom of the screen, or open the Data menu and select open file. Find the appropriate subdirectory and select your data file (for details on use of file dialog boxes, see the User Interface section).

If 3DTOM is unable to open or read the file, an error message will be issued. Check your file to make sure that it is a plain ASCII text file with the proper format. If the file is read successfully, a data information window will appear, summarizing the contents of the data file. If you wish, you can toggle between this display and a file viewer (see Menu Structure).

Options

If you wish to specify the units of measure for the data set, open the Options menu and select Units and Define. Set the appropriate buttons. The units you specify will be shown on all graphic displays. Save the options settings to a file in the same subdirectory. If you have previously saved the options settings, you can retrieve them from the saved file.

Examine Data

Inspect the data plots by pressing or clicking on F7. If you have already generated equivalent plots using the spreadsheet software, compare them to make sure the data were read correctly. Check for extreme outliers that may indicate errors in the data. View a ray-path stereoplot to evaluate the angular coverage of your data set. For ideal angular coverage, the entire plot will be filled with data. For two-dimensional data sets, ideal coverage corresponds to data spanning all possible directions in the plain. On a stereographic projection data points would extend along a line or curve all the way across or around the plot.

If there are outliers in the data plots, you can return to the original file and attempt to find the problem, or you can let 3DTOM remove the outliers by opening the Data menu and selecting Clean. Enter the acceptable range of velocities in the dialog box provided. A file dialog box will open so that you can enter a file name for saving the cleaned data. Generally it is desirable not to overwrite the original data file, so use a different name.

Check Default Grid

When you are satisfied with the data, check the default grid that 3DTOM has constructed for your data: from the Model menu select View Grid. If you want to change the grid setup (boundaries and/or numbers of nodes), create a new grid by selecting Model Generate.

Set Up Starting Model and Apply Constraints

The default starting model has a uniform velocity (or attenuation), equal to the average value calculated from the data. You can use the Model Generator to create different starting models with uniform, layered, or checkerboard structures. The Model Editor can be used to modify the starting model and to apply constraints to single nodes or groups of nodes.

If you created a borehole log file, and you want to modify the initial model to include the borehole information, open the file through the File menu. See the Files section of the User's Manual for details on including borehole information and applying related constraints.

Before running the tomographic calculations, you may wish to save the starting model for use in subsequent runs. Use the Model or File menus, or the Alt-F2 key shortcut. The default filename extension is '.3DM'; save the file in the same subdirectory as the data.

Set Inversion Parameters

Set the inversion parameters, as described in the section on Options. If you will be using the same parameters for later runs, save the options in a '.3DO' file in the data subdirectory.

Run the Inversion Routine

Begin the tomographic calculations through the Run menu or the F10 key. When the calculations are finished, a new window will appear, containing information on convergence, and the model window will be updated. You can use the F6 key to cycle among the windows.

Save the Calculated Model and Other Output

By default, the model is saved in a file called '3DTOM.3DM' in the current directory. If you don't want it to overwritten in subsequent runs, you should save it under a unique name (Alt-F2 or Model menu.)

If you wish to save other information calculated during the run (residual information, or calculated travel times or amplitudes) use the File menu to do so. This information is not saved by default.

View the Final Model.

Pressing F8 will enable you to see contoured displays of 2-D slices of the 3-D model, perpendicular to X, Y or Z, and coinciding with planes of grid nodes. If you are running under Windows, you can copy any plot to the Windows clipboard by pressing Alt-PrintScreen - you can then paste it into other Windows documents. Contour plots are described under 3D Models .

If you wish to view a contour plot on a plane with an orientation other than those available through F8, use the Model menu to select Extract (see 3D Models), and specify the plane you wish to view. You can also use Extract to determine the model values at any set of points you specify.

Compare Calculated Model With Other Models

You may want to compare the results of this run with those of previous runs (e.g., with different inversion parameters, different starting models, or different data sets from the same site). To compare models, select Compare from the Model menu (see 3D Models).

Export Model

For presentation-quality output, you can export the model in a format readable by commercial graphics software, as described in the 3D Models section.

Check Sampling Density

It is useful to view plots of raypath sampling density in order to evaluate how resolution may vary within the grid. From the Model menu, select Sampling Density. You will be able to view contoured slices showing the number of rays crossing each cell in the grid. Higher numbers in general indicate better resolution.

Recommended General Strategy

The following guidelines have been found useful for efficient processing of tomographic data. In general, it is recommended that processing progress from fast and simple to more complex and time-consuming approaches as necessary.

Initial inversions can be carried out quickly with straight-ray processing to decide whether the velocity contrasts may be high enough to warrant detailed ray-tracing. Two or more straight-ray inversions with different initial models (for example, uniform or horizontally-layered) can be helpful in evaluating the degree of nonuniqueness associated with the data set. It is also important at an early stage to check the output file containing information on residuals; large residuals may be due to incorrect travel-time picks or typographical errors, and the associated data should be re-examined. If the high-residual data appear to be correct, the discrepancies may indicate the need for curved-ray processing.

Remember that the mathematical resolving power of the data, which depends on the number and geometry of raypaths, may differ significantly from the physical resolution. Representative data traces should be analyzed to determine the dominant frequency of the first-arriving energy, which can be combined with an average velocity estimate to provide an approximate wavelength and an indication of the physical resolution possible. The Bureau software BOMSPS (Bureau of Mines Signal Processing Software; *14*) can be used for this analysis. If the default voxel size is much smaller than the wavelength (for example, less than 10%), it may be worthwhile to change to a somewhat coarser voxel grid, which will reduce execution time without sacrificing real resolution. If the default voxel dimensions are larger than a few wavelengths, a somewhat finer grid may be justified, at least locally.

For curved-ray inversions with a uniform-velocity starting model, at least one straight-ray iteration should be performed first. The reason for this is efficiency: with uniform velocities, the ray paths will be straight lines, and there is no need to carry out the ray-tracing calculations.

Constraints should be considered whenever a reconstruction is based on incomplete ray-path coverage, and is thus nonunique. Constraints may also be applied with more complete data sets in order to reduce the effects of noise and errors in the data. When a reconstruction appears to be geologically unreasonable in terms of the range or distribution of velocity, the data should first be checked for errors as described above; if the problem remains, a progression of constraints should be applied, beginning with the most objective

(generally borehole velocity measurements) and continuing to more subjective or uncertain constraints, such as maximum and minimum velocity limits, and overall structural/stratigraphic configuration. Uniformity over groups of grid nodes should always be recognized as an uncertain constraint, and an appropriate fuzziness should be applied.

FORWARD MODELING

It is useful to evaluate the resolving power of any experimental geometry by generating and inverting synthetic travel-time data. You must first set up a travel-time file containing the source and receiver locations (as described in Data Files), and either measured or "dummy" travel times.

Start 3DTOM and read in the data file. Generate a layered or checkerboard model, and save it for later comparison. Use the Options dialog box to specify only one iteration, using straight or curved rays as appropriate. Press the F10 to begin the calculations. When they are finished, open the File menu and select 'Save Calculated Times'. The file you save will contain a synthetic data set, with the same source and receiver locations as the original data file, but with calculated data replacing the measured or dummy values.

Then close the data and model files (or re-run the program), and load the synthetic data. Tomographic inversion of the synthetic data should ideally produce an accurate reconstruction of the original checkerboard or layered model. Make sure that your starting model for the tomographic inversion has the same grid geometry as the 'target' model (i.e., the model used to generate the synthetic data), but a different velocity or attenuation structure (e.g., a uniform model.) Carry out the inversion, and compare the resultant model with the target model (e.g., construct difference or ratio tomograms, and/or calculate the correlation coefficient.) The correlation coefficient will give an overall measure of the accuracy of the reconstruction, and thus an indication of the resolving power of your data set. Difference or ratio tomograms will show how the accuracy of the reconstruction, and the resolving power of the data, vary throughout the grid.

TUTORIAL EXAMPLES

Demo1 Data Set

This data set includes cross-borehole high-frequency seismic travel times for four vertical boreholes in the Colorado School of Mines Edgar Mine facility. The rock mass at this site consists primarily of gneiss intruded by granite. Laboratory measurements of P-wave velocity on intact specimens averaged 4.66 m/ms (16.4 ft/ms). The data were processed in two dimensions for each borehole pair by Jessop et al (*15*), and a geological interpretation of the results is presented in that paper. The data were re-processed in three dimensions by Jackson et al (*16*), and issues of resolution are discussed therein. The objective of the initial study was to locate zones of fractured rock prior to a simulated stope leaching experiment.

Through the File menu, choose Change Directory and select DEMO1. Open the file DEMO1.3DD. The spatial coordinates in the file are given in feet, and times in ms, so the default starting model constructed by 3DTOM will have velocities in ft/ms. From the Options menu, select Retrieve, and load the options from DEMO1.3DO, which specifies these units.

Examine the experimental geometry by selecting View Grid from the Model menu. The initial view is along the Z axis, parallel to the vertical boreholes (i.e., a map view). White points are grid nodes; yellow circles represent receiver positions, and cyan asterisks represent source locations. Raypath coverage is confined to the vertical planes between boreholes, and it is apparent that large portions of the grid are unsampled by the ray paths. Rotate the grid by pressing the up arrow until you are looking at a vertical section view, parallel to the Y axis. The source and receiver locations are restricted to particular vertical ranges; again it is apparent that many of the grid cells are unlikely to have been sampled by any of the measured rays.

The data summary dialog box will display the range and mean of the velocities calculated for the individual measurements. Note that the calculated velocities extend significantly above and below the laboratory data. Reduced velocities may be due to fracturing. The much higher velocities may be due in part to overburden stress or to errors in travel-time picking with noisy data. It is important to eliminate erroneous data as much as possible, because they can produce substantial artifacts in an ill-conditioned inversion.

Press F7 to view the data graphs. The time-distance plot shows a general trend of increasing travel time with increasing distance, with groups of data points both above and below the best-fit line, but no obvious outliers. The velocity histogram shows a more or less normal, unimodal distribution, again without obvious outliers. In the raypath stereoplot, the five vertical planes surveyed are defined by lines of data points extending radially out from the center of the plot. The color-coding shows that the lowest velocities are generally confined to one of the planes, and the highest velocities occur in a different plane. The distribution of ray-path orientations can be seen to be quite limited, with very limited azimuthal coverage, and with no steeply-inclined rays (near the center of the plot), as is typical for crosshole data sets.

It is usually instructive to carry out a rapid straight-ray inversion and examine the spatial sampling density and the residuals. Through the Options menu, you can open the Inversion Parameters dialog box and set the maximum iteration number to a relatively small value such as five, and the number of straight-ray iterations to the same value. Press F10 to begin the SIRT calculations. When they are finished, save the residual information via the File menu; you can examine the file later using any text editor. To further examine the sampling density, select Sampling Density from the Model menu. Selecting the Statistics option will cause 3DTOM to display basic statistics about the sampling density; for this data set and the default grid, it shows that only about 12% of the cells are sampled by ray paths. Selecting the Contour option leads to a contoured plot of sampling density on planes perpendicular to the X, Y or Z axis. Try viewing the plots perpendicular to Z (i.e., contoured horizontal planes), which will clearly show fairly high sampling density along and between the boreholes, and none elsewhere. With this data set, it is impossible to resolve velocity variations over a large fraction of the grid volume, but local resolution may still be adequate.

Examine the computed velocity model by pressing the F8 key. Even after only five iterations, there are very strong variations in velocity. Note that these variations are confined to the areas with nonzero sampling density; unsampled regions remain at the uniform default initial velocity, and appear in the reconstruction as featureless, uniform regions of average velocity. A prominent resolved feature is the low velocity zone extending subhorizontally at a depth of about 20 feet, and a Y coordinate of -3 feet, with velocities much lower than those measured for intact specimens in the laboratory. Additional drilling showed this region to be strongly fractured (15).

You may sometimes wish to carry out more iterations after viewing a computed model. Save the the model by pressing Alt-F3 before continuing, if you wish to save the intermediate result. Return to the Inversion Parameters dialog box, and set the maximum iteration number to one, and the number of straight-ray iterations to zero. 3DTOM will then do one curved-ray iteration, starting with the model already calculated, when you press F10. During the ray-tracing calculations, you will see the locations of each source and each receiver in turn as each ray is traced. After the calculations are finished, bring the Residuals window to the foreground by clicking on it or by pressing F6 twice. You will see that the RMS residual decreased steadily for the first five (straight-ray) iterations, and then increased somewhat as a result of the curved-ray iteration. This increase nearly always occurs for the first curved-ray iteration after several straight-ray iterations; it is caused by the straight-ray iterations making corrections in the "wrong" places. With additional curved-ray iterations, the process generally resumes converging to a solution (i.e., residuals tend to decrease).

When a data set is comprised of several interlocking crossborehole surveys, as in this example, the best resolution is in the planes between boreholes. It is therefore of interest to view contour plots of the velocites in these planes. For the DEMO1 data set, two of these planes are approximately perpendicular to the coordinate axes, and these two can thus be viewed by pressing F8 and finding slices near X = 0 and Y = -3. The other surveyed planes are oriented diagonally across the grid. To view contour plots of these planes, you must first note the coordinates of the boreholes of interest. Select Display Mode from the Data menu, and choose File View. This will open a scrollable text window displaying the data file contents. The data numbered from 98 to 217 were collected using the borehole pair with X,Y coordinates (-5.84, -2.4) and (0, 9.29), and Z coordinates between 10 and 20. To generate a plot of this rectangular area, select Extract from the Model menu, and choose Slice. You will have to specify the coordinates of three of the corners for the plot: use (-5.84,-2.4,10), (0,9.29,10), and (-5.84,-2.4,20) to produce a vertical slice between the boreholes.

To continue to the second demonstration data set, open the Data menu and select Close, and then open the File menu and use the Change Directory dialog box to log onto the DEMO2 subdirectory.

Demo2 Data Set

This data set is much larger than DEMO1, with 2365 measured P-wave travel times from a deep underground silver mine in the Coeur d'Alene district of northern Idaho. The site consists primarily of quartzites and argillites, cut by a regionally-important fault and by two ore veins. The coordinate system for the data is right-handed, with units of meters and the Z axis positive upwards. Times are in ms. The data were collected in an attempt to qualitatively evaluate stress conditions, to help reduce the risks associated with rockbursts (*17*).

Open the data and options files as described above for the DEMO1 data set, and view the grid. Source and receiver locations are within the underground mine openings on two levels. Note in the Data Information dialog box that velocities calculated for the individual measurements range from 0.4 to 10.6 m/ms. The time-distance plot shows a strong linear trend, with a few outliers below and a gradational decrease in data density above. The velocity histogram shows the same high-speed outliers, and suggests that data above 6 m/ms are erroneous; velocities taper off gradually at the low end, and no obvious outliers are present there. The ray stereoplot shows two groups of ray paths: a subhorizontal group (around the perimeter of the plot) with sources and receivers on the same level, and a steeply-

inclined group for transmission between levels. The velocities are somewhat lower for the intralevel group, and the fast outliers can be seen within the interlevel (steep) group. Overall, the angular coverage is much more extensive than it is in the DEMO1 data set.

Elimination of outliers prior to tomographic inversion can be done easily by selecting Clean from the Data menu. Enter 6 as the maximum velocity. 3DTOM will write a new data file, excluding data with higher velocities (it is best not to overwrite your original file). Open the file with the cleaned data, and re-examine the plots; the outliers will be gone.

Return to the grid viewer and rotate the grid so that the view is along the Y axis, with the Z axis upward. Note that the cells are not equidimensional - the default grid has the same number of nodes in each direction (22 in this case). Because the interlevel ray paths are generally steep, the vertical resolution between levels is probably not sufficient to warrant so many nodes in the Z direction (with limited viewing angles, velocity variations tend to "smear" in the direction of the ray paths). Reducing the number of nodes in the Z direction by half would therefore probably not have much effect on resolution, but would speed up the calculations significantly. Open the Model Menu and select Generate to set up a new grid with uniform velocities. In the grid dialog box, set the number of Z nodes to 12; in the velocity dialog box, the default is the average value for the data, but you can use different values for testing robustness, as described below.

A single iteration using straight rays will provide information on ray-path coverage: approximately 70% of the cells are unsampled, especially those around the edges and sides of the grid. Contoured sampling density plots show good coverage on the horizontal planes (perpendicular to Z) corresponding to the mine levels, and on vertical planes (perpendicular to X) between X coordinates of about -20 and +5. In these regions, one may expect reasonable resolution; outside these well-sampled areas resolution will be poor.

Resolution can be examined in more detail by forward modeling. Generate a new model, selecting the "checkerboard" velocity pattern. Save this model for later comparison, and carry out one iteration (use curved rays if you want to model travel times as accurately as possible). Save the computed "synthetic" travel times (through the File menu). Read in the synthetic data, and carry out ten to fifteen iterations. Compare the results with the original checkerboard model: calculate the correlation coefficient (which will describe the overall accuracy of the reconstruction), and view contour plots of the reconstructed velocities (which will show areas of good and poor resolution).

A final technique for evaluating the reliability of the velocity reconstructions is to perform multiple inversions using different starting models. Read in the travel times (the real data, not the synthetic). Generate an arbitrary starting model (uniform, layered, or checkerboard) and carry out the calculations to produce a tentative final model. Generate a different initial model, and compute a second tentative model. Repeat this several times, and then use the Model menu (Compare) to calculate an average of the tentative models, and standard deviations. Low standard deviations will mark regions of the grid for which the final model is insensitive to details of the starting models, and for which the final reconstruction can be regarded as reliable; high standard deviations will conversely mark regions of poor robustness and reliability.

Figures 12-14 illustrate some of these resolution and reliability tests for the DEMO2 data set, using the default grid (22x22x22 nodes) and the uncleaned data. The sampling density (Figure 12), as noted above, is highest in the two horizontal levels and in the vertical region near the X=0 plane. The checkerboard resolution test (Figure 13) shows poor resolution around the edges and sides of the grid, fair resolution through most of the interior, and very good resolution in the most densely sampled regions. The final velocity model

shown in Figure 14a is the average of eleven curved-ray reconstructions based on different uniform starting models. The standard deviations (Figure 14b) have a constant value in the unsampled areas, lower values through most of the interior, and near-zero values for the most densely sampled regions. Standard deviations higher than the background value occur in particular areas, often near the transition from unsampled to sampled parts of the grid; in these areas the reconstructions are very sensitive to the details of the starting models (for curved-ray processing) and the results are therefore unreliable in these particular areas.

TECHNICAL OVERVIEW

3DTOM takes as input a data set of travel-time measurements or reduced amplitude measurements, and computes a three-dimensional model of velocity or attenuation in the material containing the ray paths. Data may be recorded for seismic or electromagnetic waves, and the resultant tomographic images will provide information about the spatial variation in elastic or electrical properties in a mine site or other medium of investigation. This technical overview briefly summarizes the mathematical foundation of 3DTOM's imaging method, and details the techniques used by 3DTOM for ray tracing.

THE SIMULTANEOUS ITERATIVE RECONSTRUCTION TECHNIQUE

3DTOM, like its predecessors, is based on the simultaneous iterative reconstruction technique, or SIRT (*18*, *5*, *3*). Calculations are based on a three-dimensional rectilinear grid of node points, with intervening volume elements or voxels. Values of velocity or attenuation are specified at the nodes, and calculated within voxels by multiple linear interpolation (*3*).

SIRT calculations modify an arbitrary initial velocity model by repeated cycles of three steps: forward computation of model travel times, calculation of residuals, and application of velocity corrections. The cycle repeats through a number of iterations until one of several termination criteria is fulfilled.

Forward calculations (ray tracing and travel-time computation) in 3DTOM may be carried out in several ways, ranging from fast and approximate to time-consuming and accurate. The straight-ray approximation allows very rapid calculation of model travel times, but its validity diminishes with large velocity contrasts. Curved-ray calculations involve a great deal more computation, but are more accurate for strong contrasts. Ray-tracing methods are described in a subsequent section.

Travel-time and attenuation calculations both involve summation along calculated raypaths. The travel time for ray i (i = 1,N where N is the number of observations) is:

$$t_i = \sum_{j=1}^{M} p_j d_{ij} \qquad (1)$$

where d_{ij} is the distance traveled by ray i in voxel j (j = 1,M where M is the number of voxels), and p_j is the average slowness of the ray along the path segment within voxel j. Note that d_{ij} is nonzero only for voxels sampled by ray i.

Amplitude calculations similarly sum attenuation along the raypath. For a plane wave traveling in the direction of ray i, the amplitude is governed by:

$$-\ln(A_i/A_0) = \sum_{j=1}^{M} \alpha_j d_{ij} \qquad (2)$$

where A_0 is the source amplitude and α_j is the coefficient of absorbtion or attenuation for voxel j. Note the similarity between equations 1 and 2. However, with point sources of energy, there are other important factors that affect measured amplitude, including geometrical spreading, source radiation pattern (angular dependence of radiated energy), and angular sensitivity of the receiver. These factors must be accounted for by a step of data reduction prior to inverting with 3DTOM (see the section on Amplitude Inversion for more details). Assume for example that a dipole antenna radiates electromagnetic wave energy with an amplitude proportional to $\sin(\theta_T)$, where θ denotes the angle between the dipole axis and the direction of propagation, and the subscript T refers to the transmitting antenna. Also assume a dipole receiving antenna with a sensitivity proportional to $\sin(\theta_R)$. Finally assume spherical spreading of the wave energy, leading to amplitude decay according to 1/r, where r is the distance traveled by the wave from the source. Then

$$-\ln(A_i/A_0) - \ln(r) + \ln(\sin(\theta_T)) + \ln(\sin(\theta_R)) = \sum_{j=1}^{M} \alpha_j d_{ij} \qquad (3)$$

The first term on the left-hand side represents the total intrinsic attenuation (amplitude drop due to absorption) between source and receiver, the second is the geometric spreading loss, and the remaining two terms are angular correction factors for transmitter and receiver. The left-hand side of equation 3 is here referred to as the reduced amplitude; this must be calculated first and stored in a data file for input to 3DTOM, in order to perform an amplitude inversion.

The quantities on the right-hand side of equations 1 and 3 are computed by 3DTOM and compared with the left-hand quantities. The residual for each ray (difference between measured and calculated travel times or amplitudes) is used to calculate incremental correction factors for all of the voxels sampled by that ray. The incremental correction factors are accumulated for all rays before being applied to the voxels (hence the term "simultaneous" iterative reconstruction). For a particular ray path *i*, the incremental slowness correction Δp_{ij} applied to each voxel *j* traversed by that path is proportional to the travel-time residual Δt_i and to the path length d_{ij} of the ray within the voxel:

$$\Delta p_{ij} = \frac{\Delta t_i d_{ij}}{N_p \sum_{k=1}^{M} (d_{ik})^2} \qquad (4)$$

where N_p is the number of ray paths that sample voxel *j*. The incremental slowness corrections are summed to obtain the net correction for each voxel:

$$\Delta p_j = \sum_{i=1}^{N} \Delta p_{ij} \qquad (5)$$

Correction factors for each grid node are then obtained by averaging the corrections calculated for each voxel attached to that node.

RAY-TRACING METHODS

Ray-tracing calculations in 3DTOM can be carried out using four different methods: 1) approximation as straight lines; 2) a "ray-bending" approach (*12*); 3) a "network theory" or "shortest path" approach (*11*); and 4) a hybrid approach that uses both the network and raybending methods. The "ray-shooting" algorithm used in BOMCRATR (*5*) was not considered because of the shadow-zone problem and other computational aspects that would make it difficult to extend to three dimensions. Similarly, the numerical wavefront propagation approach used in MIGRATOM (*3*) was ruled out because it would be prohibitively inefficient in three dimensions. Extensive testing with two-dimensional models showed the bending and network algorithms to be faster and more stable, though less accurate, than the wavefront propagation technique. However, the hybrid approach achieves the same high accuracy as wavefront propagation, while executing more quickly.

Ray-bending Approach

The ray-bending approach (*12*) is based on iterative subdivision of an initially straight raypath, joining the endpoints (source and receiver locations), into an increasing number of conjoined (but not collinear) straight segments. Each iteration in the process consists of a "division phase" and an "adjustment phase." A brief description of the method follows.

For the initial iteration the process works as follows: a straight raypath is constructed from source to receiver, and the midpoint is calculated, dividing the original path into two segments. The velocity gradient is computed at the midpoint, and the midpoint is displaced according to the local gradient. The gradient at the new midpoint is then calculated, and the midpoint location is again adjusted. This "adjustment phase" is repeated one or more times before proceeding to the next iteration.

In the second iteration, the same process is applied to each of the segments produced in the first iteration, with a "division phase" followed by an "adjustment phase." Subsequent iterations proceed in the same manner, with the number of segments doubling each time, producing more smoothly curved paths. At each iteration, the travel time along the path is calculated, and iterations terminate when a stable minimum is attained. Figure 15 illustrates the process for a velocity model in which velocity increases linearly with depth. The iterative adjustments converge to a path that closely approximates a circular arc, as it should for a uniform velocity gradient. Five iterations result in a path containing 32 linear segments, with a calculated travel time that is usually accurate to within a few percent.

Network-theory Approach

This approach lends itself naturally to application in tomographic calculations because both are based on a fixed grid of node points. The network method involves calculation of the minimum time path along straight line segments between grid nodes. For each source point, the algorithm determines the first-arrival path to all other nodes in the grid, resulting in a "shortest-path tree", with its trunk at the source point and with branches extending to all possible receiver locations.

In its most efficient implementation (*11*) the network approach can be conceptually related to Huygens Principle (*19*), because the calculations are carried out along an

expanding "wavefront." From the source node, travel times are calculated along straight line paths to each immediately neighboring node. These nodes are analogous to the wavefront points for the Huygens' approach in that they are then treated as "sources" for the subsequent step. They differ from the Huygens' wavefront points in the fact that their position is always a pre-determined grid node location. The calculations continue until the "wavefront" traverses the entire network; in other words the algorithm calculates first-arrival travel times from one source point to each of the other points in the grid.

Figure 16 illustrates a set of ray paths computed using the network approach, with a single source (on the left) and an array of receivers (on the right). Paths were calculated for a model with a uniform vertical velocity gradient, i.e., with velocity increasing linearly with depth, and with a grid of 121 nodes (11 by 11). The computed paths are fairly crude, angular approximations to the circular arcs expected with a uniform velocity gradient (compare Figures 15 and 16).

The suite of ray-tracing approaches offers a range of choices for speed and accuracy. The straight-ray approximation is fastest, but generally least accurate; the ray-bending method is still very fast and more accurate; and the network approach takes the longest time to compute. However, a problem with ray-bending is that the process may converge to a local minimum rather than a global minimum time path. In other words, the result of the calculations may be a legitimate ray path, but one that is not the first arrival. This is illustrated in Figure 17, which shows ray paths calculated for a velocity model in which velocity varies sinusoidally with depth (resulting in alternating fast and slow horizontal layers). Paths are shown for a single source to an array of receivers. They show a strong tendency to follow the fast horizontal layers, while cutting across the slow layers at a fairly steep angle in order to minimize travel time. However, there are somewhat erratic jumps from one fast layer to another. In several cases raypaths travel upward above the receiver and then return downward to it. These are local-minimum or secondary-arrival raypaths.

Hybrid Ray Tracing

Paths calculated for the same velocity model using the network approach are shown in Figure 18. Again the paths show a strong tendency to follow the fast horizontal layers, and here there is no overshoot - the paths all correspond to global minimum travel time, first-arrival paths. They are rather angular and appear "unphysical," but they serve as good initial estimates of the minimum-time paths, and they can be iteratively improved by the bending algorithm. The results of this hybrid approach are shown in Figure 19, which shows high-accuracy paths that follow the appropriate high-velocity layers, and curve smoothly to yield global minimum-time ray paths.

The hybrid approach perfoms both network and bending calculations, and is therefore the most costly in execution time, but it is the most accurate of the approaches available.

SEISMIC IMAGING

Seismic imaging is here used broadly, including acoustic, elastic, ultrasonic, and seismic wave phenomena. You may invert amplitudes to obtain an image of seismic attenuation properties, or invert travel times to obtain a seismic velocity map.

Wave Modes

It is most common to use the first-arriving energy for determination of travel times and/or amplitudes, in part because first arrivals are the easiest to recognize. However other, later arrivals may also be used, provided that 1) they can be consistently identified and 2) that they correspond to transmission rather than reflection. 3DTOM is unable to calculate reflected raypaths, so using reflected raypaths will produce unreliable results.

First arrivals may be directly transmitted or refracted through a high-velocity zone or layer (often referred to as "head waves"). Reflected energy is never the first arrival. When first-arrival travel times are used, curved-ray processing (see raytracing; options) is usually most appropriate, although straight rays may sometimes be sufficiently accurate.

If shear (S-) waves are consistently identifiable, they are particularly useful for imaging. Because they travel with lower velocities than compressional (P-) waves, they have shorter wavelengths (for the same frequency) and thus may provide better resolution. In addition, if both compressional and shear velocity (Vp and Vs) reconstructions are calculated (using the corresponding travel times), you can produce a ratio tomogram Vp/Vs, which is useful for interpretation purposes.

Seismic Velocity and Material Properties

P-wave velocity is related to the shear modulus (or modulus of rigidity) μ, the bulk modulus (or modulus of incompressibility) K, and density ρ by:

$$V_P=\sqrt{(K+4\mu/3)/\rho} \quad (6)$$

(e.g., *20*). S-wave velocity depends only on the shear modulus μ and density ρ:

$$V_S=\sqrt{\mu/\rho} \quad (7)$$

Typical velocities and Q-factors (see below) for various geologic media under near-surface conditions are summarized in the following table, compiled from Telford et al (*21*), Dobrin (*22*), Buchanan and Jackson (*23*), and USBM data.

Table 1. Seismic Velocity and Attenuation for Various Rock Types

material	Vp, km/s	Vs, km/s	Q
air	0.3	0.0	
water	1.5	0.0	
loose sediments	0.3 - 2.3	0.5	
coal	1.2 - 3.5	0.6 - 2.0	20 - 70
shale		1.5 - 4.2	30 - 70
sandstone	1.5 - 4.2		20 - 50
limestone	2.5 - 6.0	2.9 - 3.0	65 - 650
granite	4.0 - 5.6	2.9 - 3.0	100-150
basalt	5.4 - 6.4	3.2	70
gabbro	6.2 - 6.7	3.4	
salt	4.2 - 7.0		

Although velocity varies inversely with density (for constant elastic properties), in geological materials velocity generally increases with increasing density. This is due in part to a strong increase in K and μ in denser rocks such as basalts and gabbros, and it is also due in large measure to the effects of porosity (which reduces bulk density and sharply decreases velocity).

The ratio Vp/Vs depends only on the elastic moduli μ and K, and is independent of density:

$$V_P/V_S=\sqrt{K/\mu+4/3}$$

This ratio can be computed from P- and S-velocity tomograms (see Comparing Models.) Vp/Vs is also related in a simple way to Poisson's ratio σ:

$$(V_P/V_S)^2=\frac{2-2\sigma}{1-2\sigma} \qquad (8)$$

$$\sigma=\frac{(V_P/V_S)^2-2}{2(V_P/V_S)^2-2}$$

You can compute Poisson's ratio by importing 3DTOM model files into a spreadsheet or graphics program that allows you to enter formulas such as the above. Poisson's ratio generally decreases with increasing shear modulus

$$\sigma=\frac{3K-2\mu}{6K+2\mu} \qquad (9)$$

from a maximum value of 0.5 for a perfect fluid ($\mu=0$) to typical values of about 0.4 for unconsolidated sediments, 0.33 for sedimentary rocks, and 0.25 for igneous rocks (*20*).

Attenuation and Material Properties

During the passage of a seismic pulse through a volume of rock, the volume is compressed, dilated, and/or distorted, before returning approximately to its original state. Due to imperfect anelasticity, some of the wave energy is dissipated by this stress-strain cycle. The exponential attenuation coefficient α thus depends on material properties (anelasticity), and on the number of stress-strain cycles per unit raypath length (i.e., on the inverse of wavelength). α is therefore frequency-dependent, with greater attenuation of higher-frequency waves in a given material. For this reason, attenuation is commonly described alternatively in terms of a frequency-independent, dimensionless parameter Q, called the quality factor (high Q values indicate low attenuation, and vice versa):
Q can thus be calculated using velocities and inverse attenuation coefficients (δ) calculated

$$Q=\pi/(\lambda\alpha)=\omega/(2V\alpha)=\omega\delta/(2V) \quad (10)$$

by 3DTOM, using self-consistent units for V, δ, and ω (e.g., m/s, m/Np, and Hz).

Factors Affecting Seismic Velocity and Amplitude

In addition to the lithologic variation described above, seismic velocities are also sensitive to pressure, temperature, fluid saturation, fracturing, and other factors. Detailed discussion of these effects is beyond the scope of this document; however some simple generalizations follow. Often velocity and Q are affected in the same sense as one of these factors is varied, i.e., both increase or both decrease, despite the inverse dependence of Q on V.

Velocity and Q tend to increase with increasing pressure, primarily due to closure of microcracks and porosity. Elevated temperatures decrease V and Q, mainly by reducing rigidity. For P-waves in porous unsaturated sediments, V and Q increase with increasing saturation. Fracturing reduces both V and Q.

ELECTROMAGNETIC IMAGING

Electromagnetic (EM) wave energy can be used in exactly the same way as seismic wave energy for travel-time and amplitude inversions. However, the physical properties governing EM waves are different from those controlling seismic waves, and different information is thus obtained by EM and seismic imaging.

The behavior of EM waves varies with distance R from the source, and three different zones are generally defined in terms of the EM wavelength λ (e.g., *24*): the *near field zone* ($R << \lambda$), the *intermediate (induction) zone* ($R \approx \lambda$), and *the far-field (radiative) zone* ($R >> \lambda$). EM tomography is generally carried out with source and receiver separated by many wavelengths, so that most of the survey area is in the far field.

EM Waves and Material Properties

For EM waves, the wave velocity V and the attenuation coefficient α are related to material properties and source frequency by the following relationships (*24*):

$$V=\frac{1}{\sqrt{\frac{\mu\epsilon}{2}}\sqrt{\sqrt{1+\frac{1}{Q^2}}+1}} \quad (11)$$

$$\alpha=\omega\sqrt{\frac{\mu\epsilon}{2}}\sqrt{\sqrt{1+\frac{1}{Q^2}}-1} \quad (12)$$

where ω is angular frequency, μ is magnetic permeability, ϵ is permittivity, and Q is called the quality factor. 3DTOM model files contain the inverse of α, termed the attenuation distance or "skin depth" δ. For linear, isotropic materials, the quantities in these equations are related to other properties in the SI system as follows:

$$\begin{aligned} \mu &= \kappa_m \mu_0 \\ \kappa_m &= 1+\chi_m \\ \chi_m &= M/H \\ \epsilon &= \kappa_e \epsilon_0 \\ \kappa_e &= 1+\chi_e \\ \chi_e &= P/E \\ Q &= \omega\epsilon/\sigma \end{aligned} \quad (13)$$

where κ_m is relative permeability, μ_0 is the permeability of free space, χ_m is magnetic susceptibility, M is magnetization, H is magnetic field, κ_e is relative permittivity or dielectric constant, ϵ_0 is the permittivity of free space, χ_e is electric susceptibility, P is polarization, E is electric field, and σ is electrical conductivity. The inverse of Q ($Q^{-1}=\sigma/\omega\epsilon$) is often referred to as the "loss tangent."

For almost all rocks, $\chi_m << 1$ and $\mu \approx \mu_0$. Velocity and attenuation are thus controlled primarily by dielectric properties, conductivity, and frequency, and the Q parameter describes the relative importance of these. Conductivities of earth materials vary over many orders of magnitude, e.g., 10^8 (ohm-m)$^{-1}$ for native copper to 10^{-14} for quartz;). The dielectric constant at radio frequencies is 1 in air, 5-10 in many rocks, and 80 in water *(21, pp. 451-454)*. Q values much greater than one indicate that dielectric properties are important, whereas for $Q << 1$ conductivity dominates. Q increases with frequency, so at higher frequencies the dielectric properties become more important: for most rocks above 1 GHz, α and V are quite sensitive to variations in ϵ, whereas below 1 kHz they are determined entirely by σ. When conductivity and frequency are such that Q is large, the material is often referred to as "a dielectric" (at that frequency); conversely when Q is small, the material is called "a conductor".

EM Wave Velocity

In the high-Q (high-frequency/low-conductivity) regime ($Q \geq 10$), velocity is essentially independent of conductivity, and is controlled primarily by dielectric properties (note that for $\sigma = 0$, V reduces to $(\mu\epsilon)^{-1/2}$). For more conductive rocks or lower frequencies ($Q < 0.1$), velocity is virtually independent of ϵ, and is controlled primarily by conductivity, being approximately $\omega^{1/2}(\sigma\mu/2)^{-1/2}$.

Figure 20 illustrates the dependence of EM wave velocity on conductivity and on dielectric constant ($\kappa = \epsilon/\epsilon_0$), for a frequency of 60 MHz. For conductivities less than about 10^{-3} S/m (high Q), V is independent of σ and is determined entirely by the dielectric constant of the material., ranging from 3×10^8 m/s for $\kappa = 1$ (air) down to about 3×10^7 m/s for $\kappa = 80$ (water). For conductivities greater than 1 S/m (low Q), velocity is independent of κ and is determined solely by conductivity, with values less than 2.5×10^7 m/s. There is a transitional range of conductivities for which Q (at 60 MHz) is of order one, and in which velocity depends upon both σ and κ.

In Figure 21, velocity is shown for various frequencies as a function of conductivity, for a fixed dielectric constant of 8. The transition from high Q (velocity determined by κ) to low Q (velocity determined by σ) occurs at lower conductivities for lower frequencies. For measurements made at 1 GHz, for example, velocity is dielectrically controlled for conductivities as high as 0.1 S/m, whereas at 1 kHz, the high-Q behavior is confined to $\sigma < 10^{-7}$.

EM Wave Attenuation

Intrinsic attenuation is ultimately due to electrical conductivity. Note that if $\sigma = 0$ then $\alpha = 0$. For "dielectric" materials (with low but nonzero conductivity, such that $Q >> 1$ at the measurement frequency), the attenuation coefficient is (to a good approximation) proportional to conductivity and to the inverse square root of permittivity: $\alpha \approx \sigma/2(\mu/\epsilon)^{1/2}$. For low Q, attenuation is approximately proportional to the square root of conductivity, and independent of permittivity: $\alpha \approx (\mu\sigma\ \omega/2)^{1/2}$.

Figure 22 illustrates the dependence of attenuation coefficient on conductivity and dielectric constant, for a frequency of 60 MHz. In the high-Q regime (conductivity less than about 10^{-3}), attenuation is weakly dependent on κ, and it varies linearly with σ. For low Q, α is independent of κ and varies according to $\sigma^{1/2}$.

For a fixed dielectric constant of 8 and various frequencies, attenuation varies as shown in Figure 23. Attenuation for a particular conductivity increases with increasing frequency until Q becomes larger than one; further increases in frequency do not increase the attenuation coefficient.

Note that for low frequencies and/or high conductivities (low Q), $\alpha \approx \omega/V$; that is, attenuation is proportional to slowness (1/velocity.) In the low-Q regime, therefore, slowness and attenuation tomograms should be essentially identical (except for the factor of ω), since both have virtually the same dependence on conductivity. In the high-Q reqime, in contrast, attenuation is proportional to conductivity whereas slowness depends only on dielectric properties, and therefore velocity and attenuation tomograms may be significantly different.

Combining EM Velocities and Attenuations

As described above, at sufficiently high or low frequencies, EM tomographic reconstructions can be interpreted directly in terms of dielectric constant or conductivity. However for frequencies in the kHz to 100 MHz range, Q may be of order one (depending on the conductivity), and velocity and attenuation each depend on both conductivity and permittivity. Because they depend on these variables in different ways, they can be separated mathematically. When both velocity and attenuation tomograms have been calculated, they can be combined to produce maps of electrical conductivity and dielectric constant:

$$\sigma = \frac{2\alpha}{\mu V} = \frac{2}{\mu V \delta} \quad (14)$$

This requires assumption of a value for μ, and for most rocks the approximation $\mu \approx \mu_0$ is adequate. With units of m/Np for δ and m/s for V, equation 15 will give conductivity in

$$\kappa_e = \epsilon/\epsilon_0 = \frac{\omega^2 - V^2\alpha^2}{V^2\omega^2\mu\epsilon_0} = \frac{\delta^2\omega^2 - V^2}{\delta^2 V^2\omega^2\mu\epsilon_0} \quad (15)$$

(ohm-m)$^{-1}$, or equivalently, mhos/m or Siemens/m. Equation 16 gives the dimensionless dielectric constant for ω in Hz (or s^{-1}), δ in m/Np, V in m/s, and using a value of 1.11×10^{-17} s^2/m^2 for the product $\mu_0\epsilon_0$. Note however that errors or uncertainties in both the reconstructed velocities and the reconstructed attenuation coefficients will propagate into the conductivity and dielectric reconstructions, and these latter will therefore have larger error margins.

Conductivity tomograms are useful for mapping lithology (especially for outlining metallic ore bodies) and for locating conductive fluids (e.g., in-situ leaching fluids). Dielectric reconstructions are useful for distinguishing between air, dry rock, wet rock, and water, and therefore they have application in locating abandoned mines or other underground openings. Dielectric imaging also is used for location of certain contaminants.

AMPLITUDE INVERSION

Amplitude inversion involves computation of the attenuation distribution from measured amplitude data. Amplitude inversion is somewhat more complicated than travel-time inversion, for two reasons: there is an initial step of data reduction that must be carried out before amplitude data can be processed by 3DTOM, and there are additional restrictions on processing, as well as additional complexities in interpretation of results.

Amplitude computations are based on the equation

which states that the logarithmic amplitude decrease (or total attenuation) from source S to receiver R is the line integral of the attenuation rate along the path from S to R. A and

$$-\ln(A/A0) = \int_S^R \alpha \, dl \quad (16)$$

A0 refer to the amplitudes at the source and the receiver, respectively. This equation is analogous to that for calculating travel time on the basis of slowness (1/velocity); thus 3DTOM is able to calculate attenuation coefficients in exactly the same way that it calculates slowness, provided the amplitude data have undergone proper data reduction.

Because 3DTOM is most commonly used for travel-time rather than amplitude inversion, it has certain features that are convenient for velocity imaging, but not necessarily for attenuation imaging. For example, although the calculations involve slowness, by default 3DTOM displays its inverse, namely velocity. In the same way, attenuation calculations involve the attenuation rate, but by default 3DTOM displays its inverse, namely skin depth or attenuation distance. The default displays can be changed by selecting Display Mode from the Model menu, and choosing an Inverted display.

Model files created by 3DTOM will contain attenuation distance, with units determined by those of the data. For example, with locations specified in meters, and reduced amplitudes in decibels, the model values will have units of m/dB. Physically this refers to the distance a wave travels before being attenuated by 1 dB (i.e., before its amplitude is decreased to $10^{-1/20}$ (roughly 90%) of its initial amplitude.)

Ray Tracing in Amplitude Inversions

Curved ray processing should generally not be used for amplitude inversion because raypath curvature is controlled by the velocity distribution, not the attenuation distribution. When working with amplitude data, the values in any model refer to attenuation rather than velocity, and it is incorrect to calculate ray paths based on an attenuation model. An exception to this rule is EM amplitude inversion in the low-Q regime, where slowness is proportional to attenuation; curved rays may therefore be used in that case.

Interpretation of Attenuation Tomograms

The usual pitfalls related to nonuniqueness apply equally to velocity and attenuation imaging. Here we will focus on considerations specific to attenuation images.

Perhaps the most important consideration is the inability to distinguish between scattering effects and intrinsic attenuation. This can have various surprising consequences. For example, consider an air-filled subsurface void, such as a tunnel. For EM waves, the intrinsic attenuation in air is much less than that in rock, and one might therefore expect larger amplitudes for rays passing through the void. However, the strong contrast in physical properties across the void boundaries results in very strong scattering and a sharp decrease in the amplitude of the transmitted wave. Depending on various factors including the size of the void, the balance between scattering losses and lower attenuation in air may result in positive or negative amplitude anomalies, or even no anomaly at all. In addition, interference patterns may arise from scattering off different parts of an anomaly, resulting in amplitude maxima and minima that are not correctly computed by the ray-based methods of 3DTOM.

An additional potential problem is the inability to use curved-ray processing. In a medium with sufficiently strong variation in physical properties, wave energy can be significantly focussed just as light through a lens, resulting in strong amplitude anomalies. Such anomalies will be incorrectly translated by 3DTOM into regions with anomalously high or low intrinsic attenuation in a tomographic reconstruction.

What Is Amplitude?

Elastic Wave Amplitudes.

The elastic wave equation, derived from stress-strain relations for elastic solids, describes how *displacements* propagate in space and time (*20*). Consider a point in the rock mass with position specified by coordinates (X, Y, Z). As an elastic wave passes, the rock is locally compressed or distorted, and as a result the point at time t may have moved to location (X + dX, Y + dY, Z + dZ). The displacement vector at time t is then (dX, dY, dZ); an instrument that detects displacements in the X direction, for example, will measure a value dX at time t. Continuous recording produces a time series dX(t). The total displacement at any time has a magnitude $(dX^2 + dY^2 + dZ^2)^{½}$, and to determine this you need three perpendicular single-axis sensors or a triaxial sensor.

Many types of sensors are designed to measure not displacement, but one of its time derivatives, velocity or acceleration. Engineering geophones typically produce a voltage proportional to ground velocity (not to be confused with wave velocity; the former is the speed with which the ground moves at the measurement point, and the latter is the speed with which the disturbance propagates through the rock). Other sensors, e.g., hydrophones, use piezoelectric materials to produce a voltage proportional to pressure. Combining data from different types of sensors for amplitude inversion requires careful cross-calibration and is generally not recommended.

Electromagnetic Wave Amplitudes.

The EM wave equation, derived from Maxwell's equations, describes how components of the electric field E and magnetic induction B vary in space and time (*24*). Sensors are typically electric dipole antennas (in which a current is generated, proportional to the parallel component of E) or magnetic dipole (loop) antennas (in which a current is induced, proportional to the perpendicular component of dB/dt). Measurement of the total field E or dB/dt again requires three orthogonal sensors.

Radiation.

Wave behavior for both seismic and EM sources is rather complex in the immediate vicinity of the source (*20,24*). The wavelength of the energy is often used to distinguish between "near-field" (much less than one wavelength from the source), "far-field" (many wavelengths distant from the source), and intermediate zones. For EM waves, the far-field and intermediate zones are often referred to as the radiation zone and the induction zone, respectively. For geophysical tomography, measurements are most often made in the far-field, and the calculations of 3DTOM assume this to be the case.

Wave amplitudes may vary with direction for certain types of sources. For electric dipoles (straight antennas) the magnitude of the E and B fields produced (in the far-field) is proportional to the sine of the angle between the dipole and the propagation direction (i.e., amplitude is maximum perpendicular to the antenna and minimum parallel to it.) Simila-ly, for magnetic dipoles (loop antennas, with the magnetic dipole perpendicular to the loop),

amplitude is proportional to the sine of the angle with respect to the dipole. Seismic sources may have a similar dipole radiation pattern, or a monopolar radiation (with amplitude independent of direction), or some other dependence. Explosive sources may be nearly monopolar. Directed sources (such as hammer impacts, weight drops, etc.) typically have a dipole-like radiation pattern for shear waves (with amplitude proportional to the sine of the angle between impact direction and propagation direction) and a complementary compressional wave radiation pattern (amplitude proportional to the cosine of the angle.) Continuous (or quasi-continuous) sources such as continuous coal miners, tunnel boring equipment, or drilling rigs can be used for collecting amplitude data, but it is not always obvious what the radiation pattern of such sources may be, and care is required.

What Is Attenuation?

The E and B fields of an EM wave, and the displacements of an elastic wave, decrease in magnitude as the wave travels. This decrease, or attenuation, results from three principal causes.

Geometric spreading.

Waves emanating from a point source expand spherically (in a uniform medium) with time. Neglecting for now absorption and scattering, the total energy of a spherical wavefront remains constant as the wave expands, and the energy per unit area diminishes accordingly, following the well-known inverse-square law. Energy is generally proportional to the square of amplitude (displacement, E, B, or their time derivatives); consequently amplitude decays as 1/r (where r is the distance from the point source).

For an infinitely long line source, waves expand cylindrically rather than spherically. In this situation, energy (per unit area on the wave front) decreases as 1/r and amplitudes as $r^{-1/2}$. This holds approximately for line sources of finite length, if r is much less than the source length. Conversely, if r is much greater than the source length, the source "looks like" a point dipole from the receiver, and amplitudes diminish approximately as 1/r. In a layered medium, waves may be "trapped" in a particular layer (a "waveguide"), and cylindrical spreading may be a good approximation.

Absorption (intrinsic attenuation).

In an electrically-conductive material, EM waves induce eddy currents, and consequently some of the wave energy is dissipated or converted into heat. In an imperfectly elastic material, seismic waves lose some energy to internal friction. Both of these types of attenuation are determined by the physical properties of the medium (in contrast to geometric spreading, which is determined principally by experimental geometry). Thus it is this intrinsic attenuation that is of chief interest for geophysical imaging. In a uniform attenuating medium, intrinsic attenuation causes an exponential decay of amplitude with distance, i.e., amplitude is proportional to $\exp(-\alpha r)$, where α is referred to as the attenuation coefficient. α has dimensions of inverse length, and units are typically nepers/m (Np/m) or decibels/m (dB/m). The reciprocal of α is referred to as the skin depth or the attenuation distance δ, representing the distance over which amplitude decreases to a fixed fraction of its initial value.

Scattering.

In a medium with discontinuities (e.g., layering or other sharp lithologic boundaries) wave energy is partitioned at each interface encountered by a wave, into reflected and transmitted fractions (*20, 24*). For elastic waves, moreover, this may involve mode conversion; for example a compressional wave incident on a boundary may give rise to compressional and shear waves, both transmitted and reflected (*20*). The technique of diffraction tomography attempts to utilize scattered wave energy to image discontinuities (*25*). 3DTOM however is based solely on transmission calculations, so that energy loss due to scattering is not distinguished from intrinsic attenuation.

Quantifying Amplitudes

Far-field wave amplitudes can be measured for either pulsed or continuous energy sources and used for tomographic inversion.

Pulsed systems.

These produce wave energy in short bursts; examples include hammers, explosives, air guns, sparkers, and pulsed radar systems. Measurements are made in the time-domain; that is, ground motion or EM fields are measured as a function of time, and the arrival of a wave front is recorded as one or more peaks and troughs in the time-series. Amplitude is usually characterized in terms of peak height, trough depth, or a sum of these. Recorded amplitudes generally depend on factors such as detector sensitivity and amplification, and these calibration factors must be known to determine actual ground motion.

Continuous systems.

These generate continuous or quasi-continuous wave trains; examples include continuous-wave radio transmitters, drilling rigs, and mining machines. Measurements are made in the frequency domain; that is, amplitude is determined as a function of frequency (alternatively measurements may be made in the time domain and transformed to the frequency domain.) Amplitudes are then characterized in terms of peaks in the amplitude spectrum, or by integration over all or part of the spectrum. Again the amplitudes thus determined depend in part on the characteristics of the measurement and recording system.

Linear and logarithmic scales.

Because of the exponential nature of intrinsic attenuation, it is often convenient to characterize amplitudes on a logarithmic scale rather than a linear one (in fact, as described later, 3DTOM requires amplitude data to be scaled logarithmically.) Logarithmic scaling naturally involves dimensionless ratios, or relative quantities, rather than absolute quantities; amplitudes are thus normalized to some quantity, and then the logarithm is taken of the ratio. Amplitude ratios are most often expressed in decibels (dB), defined as 20 times the decimal logarithm of the amplitude ratio (thus 20 dB corresponds to an amplitude ratio of 10). Alternative units are nepers, defined as the natural logarithm of the amplitude ratio (1 dB = 0.1151 Np).

Some measurement and recording systems report amplitudes in dB (relative to some reference amplitude.) For example, a continuous-wave (frequency-domain) radio receiver using electric dipole antennas may display amplitudes in dB relative to one millivolt. If the potential difference due to the electric field at the receiving antenna is 10

millivolts, the value displayed will be 20log10(10/1), or +20 dB. If the receiver displays a value of -10 dB, the field is $10^{[-10/20]}$, or 0.316 mV.

A common practice is to place one or more receivers close to the source, and to normalize all other amplitudes at greater distances to this "source amplitude" A0. If the linear amplitude measured by a particular receiver is A, the ratio A/A0 corresponds to the amplitude reduction along the propagation path, or the total attenuation (due to spreading, absorption, and scattering), and this can be expressed in decibels as dA = 20log10(A/A0) = 20log10(A)-20log10(A0). A value of zero indicates no attenuation (A = A0); greater negative values correspond to greater total attenuation. For a plane wave (with no spreading) the total attenuation is equal to the intrinsic attenuation rate (dB/m) times the distance traveled. This is analogous to the travel-time equation: total travel time is equal to the wave slowness (1/velocity) times distance.

Amplitude Data Reduction

3DTOM performs its calculations in exactly the same way for both amplitude data and travel-time data. Therefore it is necessary to reduce the amplitude data to a standard form appropriate for the calculations. For travel times, 3DTOM calculates the integral of wave slowness along each ray path. If, instead of slowness, we integrate the attenuation coefficient α along each ray path, the result is a calculated logarithmic amplitude drop (or total attenuation) dA. More specifically, it is the logarithmic amplitude drop dA_A due to intrinsic attenuation, or absorption (as opposed to geometric spreading and scattering.) This, then, is what we must provide as data for 3DTOM. In order to calculate dA_A, we must correct dA for the loss due to geometric spreading dA_G: $dA_A = dA - dA_G$. Because we cannot correct for scattering loss dA_S, any losses due to scattering will show up in the tomographic image as absorption. The following is a step-by-step description of the corrections that may be required.

1. Determine "raw" amplitudes. For time-domain records with pulse sources, determine the peak height, trough depth, or combination of these for the arrival of interest. Although the first arrival is usually the easiest to identify, it may not be the most appropriate for a straight-ray amplitude inversion (considerations pertaining to tomographic inversion are discussed in Amplitude Inversion). For example, if you are able to identify both refracted first arrivals and later direct arrivals, the direct arrivals are more appropriate to use because they correspond more closely to the straight paths that 3DTOM will use in the calculations. However, you should attempt to avoid using reflected arrivals, because 3DTOM will not be able to determine the correct path for calculation.

For continuous-source frequency-domain data, different wave modes cannot be identified directly by their separation in time, but you may be able to identify particular frequency ranges to use for amplitude determination. For example if you are using a drill bit or mining machine as a source, you may be able to isolate part of the spectrum where the desired signal is strong and other mechanical and electrical noise sources are minimized. However, it is unlikely that you will be able to separate direct from refracted or reflected energy. Determine the peak amplitude for the selected frequency range, or integrate over the frequency range to quantify the amplitude, for each source-receiver combination.

For either time or frequency domain determinations, you will probably first obtain a raw measure of amplitude expressed in volts, reflecting the receiver sensitivity,

amplification, and wave amplitude. If different system parameters were used for different ray paths, you will need to calibrate the data.

2. Calibrate amplitudes. Seismic receivers generally produce a voltage proportional to displacement, velocity, or acceleration, and the constant of proportionality is termed the sensitivity (for example an accelerometer may have a sensitivity of one volt per g, where g is the average gravitational acceleration at the Earth's surface, 9.8 m/s^2). It is not unusual with multichannel recording systems to use receivers with different sensitivities simultaneously, for example at different distances from the source. In such a situation, you must divide each raw amplitude by the corresponding sensitivity. Similarly, you may have different gain settings or amplification for different ray paths, and you must correct for this. Gain is often expressed logarithmically, in terms of dB (a gain of 20 dB, for example, indicates that measured values have been multiplied by a factor of ten).

3. Make spreading correction. Ordinarily it is safe to treat transmitters as point sources (provided the source-receiver distance is much greater than the source dimensions), and to treat spreading as spherical. If, however, the point-source approximation is not valid, or if wave energy is confined to a wave-guide, a cylindrical spreading correction may be more appropriate.

For spherical spreading, amplitudes decay as 1/r; for cylindrical spreading the rate is $r^{-1/2}$. To compensate for this loss, multiply each linear amplitude by the corresponding source-receiver distance (or by the square root of the distance for cylindrical spreading). (This multiplication adds a factor of length to the units of amplitude, but this will be normalized out later.) The result is an estimate of what the amplitude would be for a plane (non-spreading) wave for each ray path.

If your data are logarithmic (dB or Np), convert the appropriate distance term (r or square root of r) to the same log units (e.g., 20log10(r)), and add to the measured amplitude for each ray path.

4. Make radiation correction. If your energy source has a radiation pattern that is not monopolar (i.e., if it generates waves with amplitudes that vary according to direction of propagation), you need to correct for this. For electric or magnetic dipole antennas, for example, the radiation is often proportional to $\sin(\theta)$, where θ is the angle between the dipole axis and the line from the source to the receiver. (Tuned antennas, such as a half-wavelength antenna, have a stronger dependence on angle.) Shear waves generated by impact sources (e.g., hammer or weight-drop) also tend to have amplitudes proportional to $\sin(\theta)$, where the "dipole" axis coincides with the impact direction; compressional waves typically have a $\cos(\theta)$ amplitude dependence. For other types of sources, such as mining machinery, it may not be apparent what sort of radiation to expect, and it may be necessary to determine it by experiment. Whatever the appropriate function $F(\theta)$ is determined to be, it must be compensated for by calculating θ for each source-receiver combination, and dividing the corresponding amplitude by $F(\theta)$. This will increase the amplitude for unfavorably oriented rays, and not change amplitudes for favorably oriented rays, so that the corrected values are close to what would have been determined for a monopolar source.

With logarithmic amplitude data, convert the angular correction factor to the same log units (dB or Np). Because sine and cosine are always less than or equal to 1, their logarithms will be negative. Subtract this negative number from its corresponding

amplitude, in effect adding a positive correction.

5. Make receiver sensitivity correction. Often receivers measure only a single (Cartesian) component of the amplitude; for example many geophones measure only the vertical component of ground motion, and an electric dipole antenna measures only the component of electric field parallel to the antenna. The full amplitude can be determined from triaxial data, but must be calculated for single-axis data, by assuming that the propagation path is approximately a straight line from the source to the receiver. Further assuming the nature of the wave is known, e.g., compressional or shear, we can determine the expected orientation of particle displacement, velocity or acceleration, or E and B fields. Calculate the angle θ between this orientation and that of the receiver axis for each ray path, and divide the measured amplitude by $\sin(\theta)$. This will provide an estimate of the amplitude that would have been measured by an optimally-oriented single-axis receiver at the same location.

If you have logarithmic rather than linear amplitudes, handle the correction as described above for source radiation.

6. Compute logarithmic relative amplitude. If your corrected data are linear amplitudes A (e.g., displacements, velocities, voltages, etc), and you have measured (and corrected) the source amplitude A0 using a receiver placed near the source, calculate the amplitude ratio A/A0, and convert to either nepers (by taking the natural log) or dB (by taking the decimal log and multiplying by 20). The result should be a negative number in all cases.

If you have linear amplitude data without a direct estimate of A0, you can estimate it indirectly. Begin by converting the corrected linear amplitudes to Np or dB. Plot them as a function of source-receiver distance; generally you will see something like a linear trend with a negative slope. Extrapolating this trend backwards to zero distance will give you an estimate of A0 (in the same Np or dB units). Subtract this value from all of your logarithmic amplitudes; the result should be a negative number in all cases.

If you have logarithmic amplitude data to begin with (i.e., if your systems records amplitudes in dB relative to some arbitrary reference value), and you have a measured value for A0, subtract this from all other data to obtain the logarithmic amplitude ratios. If your data are logarithmic and you do not have a measurement for A0, determine it indirectly by extrapolation as described above. Again the result should be a negative number for all rays.

7. Change sign. The logarithmic relative amplitudes obtained by the steps above will all be negative, reflecting the fact that measured amplitudes are all less than the source amplitude, due to attenuation. To put the data into a form analogous to travel-time data, simply change the sign to positive; the values now indicate the total attenuation (or amplitude drop) along each ray path. 3DTOM will calculate corresponding values as line integrals of attenuation rate (dB/m or Np/m) along the ray paths.

SUMMARY

Three-dimensional tomography can provide a more realistic and reliable reconstruction than two-dimensional because the bending of rays out of the plane between

boreholes can be considered, and provides more flexibility because source and receiver positions do not need to be in the same plane. However, it presents special challenges, such as finding a mathematical method that is computationally efficient, but that is not prone to causing artifacts by finding local minima of travel times. The program 3DTOM provides a hybrid method that is a combination of a rapid net approach to obtain a first approximation to the path, and then a ray bending method to refine the path. It also provides a variety of methods to examine data and to display the results. The authors hope that this last tomographic computer program from the USBM will be a useful tool for geophysicists in many types of investigations.

REFERENCES

1. Worthington, M. H. An Introduction to Geophysical Tomography. First Break, v. 2, 1984, pp. 20-26.

2. Nolet, G. (ed.). Seismic Tomography, With Applications in Global Seismology and Exploration Geophysics. Reidel, Boston, 1987, 336 pp.

3. Jackson, M. J., and D. R. Tweeton, MIGRATOM - Geophysical Tomography Using Wavefront Migration and Fuzzy Constraints, USBM RI 9497, 1994, 35 pp.

4. Tweeton, D. R. A Tomographic Computer Program with Constraints To Improve Reconstructions for Monitoring In Situ Mining Leachate, USBM RI 9159, 1988, 70 pp.

5. Tweeton, D. R., M. J. Jackson, and K. S. Roessler. BOMCRATR - A Curved Ray Tomographic Computer Program for Geophysical Applications. USBM RI 9411, 1992, 39 pp.

6. Zhou, H.-W. Travel-Time Tomography With a Spatial Coherence Filter. Geophysics, v. 58, 1993, pp. 720-726.

7. Davis, J. C. Statistics and Data Analysis in Geology. John Wiley and Sons, New York, 1973, 550 pp.

8. Press, W. H., B. P. Flannery, S. A. Teukolsky, W. T. Vetterling. Numerical Recipes: The Art of Scientific Computing. Cambridge Univ. Press, New York, 1986, 818 pp.

9. Bregman, N. D. Tomographic Inversion of Crosshole Seismic Data. Ph.D. Thesis, Univ. Toronto, Toronto, 1986, 180 pp.

10. Menke, W. The Resolving Power of Cross-borehole Tomography. Geophys. Res. Lett., v. 11, 1984, pp. 105-108.

11. Moser, T. J. Shortest Path Calculation of Seismic Rays. Geophysics v. 56, 1991, pp. 59-67.

12. Um, J. and C. Thurber. A Fast Algorithm for Two-Point Seismic Ray Tracing. Bull. Seismol. Soc. Am., v. 77, 1987, pp. 972-986.

13. Singh, R. P. and Y. P. Singh. Geophysics, v. 56, 1991, pp. 1215-1227.

14. Friedel, M. J. BOMSPS: Bureau of Mines Signal Processing Software - Concepts, Expressions, and Tutorial. USBM IC 9242, 1990, 40 pp.

15. Jessop, J. A., D. L. Boreck, M. J. Jackson, D. R. Tweeton, and M. J. Friedel. Evaluation of Stope Leaching Site Using Geotomography. Paper in In Situ Recovery of Minerals II (Santa Barbara, CA, Oct. 25-30, 1992). Engineering Foundation, New York, 1994, pp. 599-616.

16. Jackson, M. J., M. J. Friedel, D. R. Tweeton, D. F. Scott, and T. Williams. Three-dimensional Imaging of Underground Mine Structures Using Seismic Tomography. Paper in Proceedings of the Symposium on the Application of Geophysics to Engineering and Environmental Problems (SAGEEP95) (Orlando, FL, Apr. 23-26, 1995) pp. 221-230.

17. Friedel, M. J., M. J. Jackson, D. F. Scott, T. J. Williams, and M. S. Olson. 3D Tomographic Imaging of Anomalous Conditions in a Deep Silver Mine. J. Appl. Geophysics, v. 34, No. 1, 1996, pp. 1-21.

18. Peterson, J. E., B. N. P. Paulson, and T. V. McEvilly. Applications of Algebraic Reconstruction Techniques to Crosshole Seismic Data. Geophysics, v. 50, 1985, pp. 1566-1580.

19. Saito, H. Traveltimes and Raypaths of First Arrival Seismic Waves: Computation Method Based on Huygens' Principle. Paper in Society of Exploration Geophysicists 59th Annual Meeting, Expanded Abstracts. SEG, 1989, pp. 244-247.

20. Richter, C. F. Elementary Seismology. W. H. Freeman and Co., San Francisco, 1958, 768 pp.

21. Telford, W. M., L. P. Geldart, R. I. Sheriff, and D. A. Keys. Applied Geophysics. Cambridge Univ. Press, New York, 1976, 860 pp.

22. Dobrin, M. B. Introduction to Geophysical Prospecting. McGraw Hill, New York, 1976, pp. 25-57.

23. Buchanan, D. J., and L. J. Jackson. Coal Geophysics. Geophysics reprint series #6, SEG, 1986, 466 pp.

24. Wangsness, R. K. Electromagnetic Fields. John Wiley and sons, New York, 1979, 576 pp.

25. Devaney, A. J. Geophysical Diffraction Tomography. IEEE Trans. Geosci. and Remote Sensing, v. GE-22, 1984, pp. 3-13.

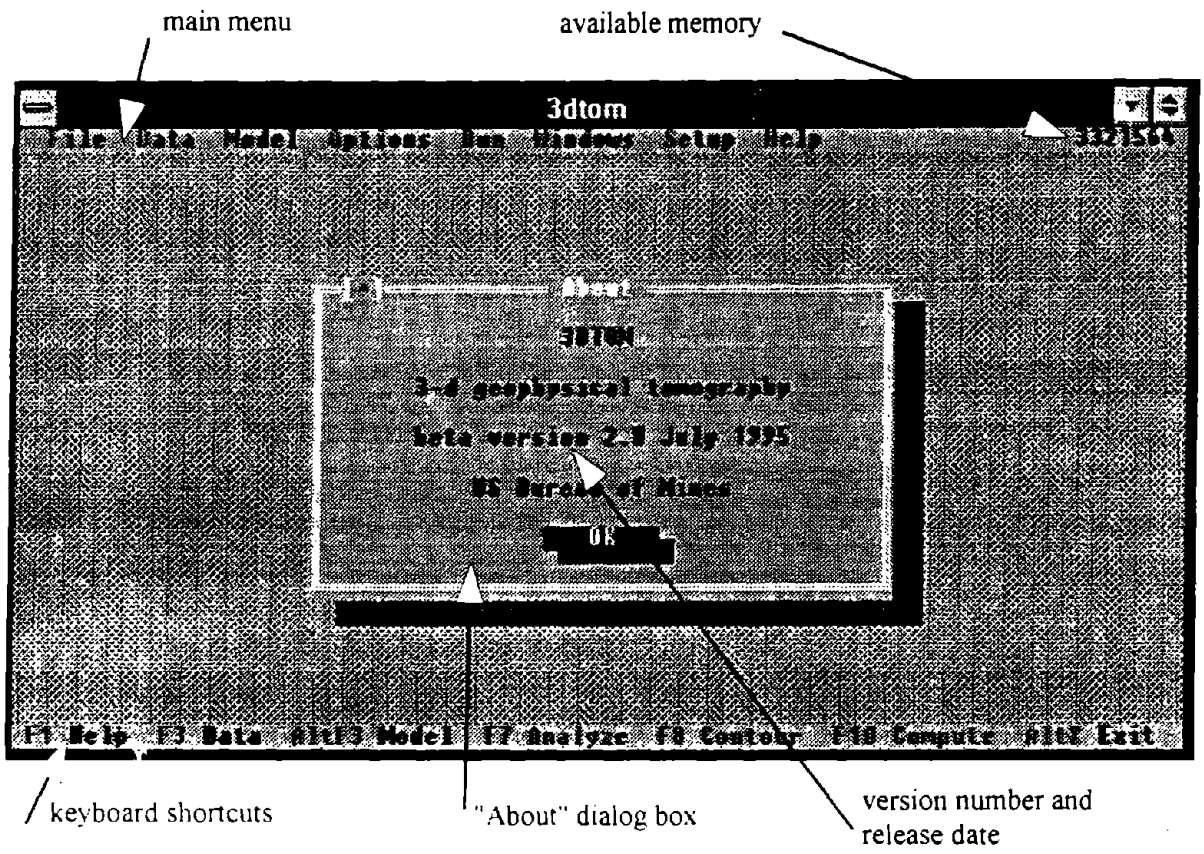

Figure 1. Introductory and main menu screen for 3DTOM (as it appears when running through Windows). Menu items may be selected by using a mouse or by pressing the Alt key together with the highlighted letter in the menu item. The shortcuts on the bottom of the screen work in the same way.

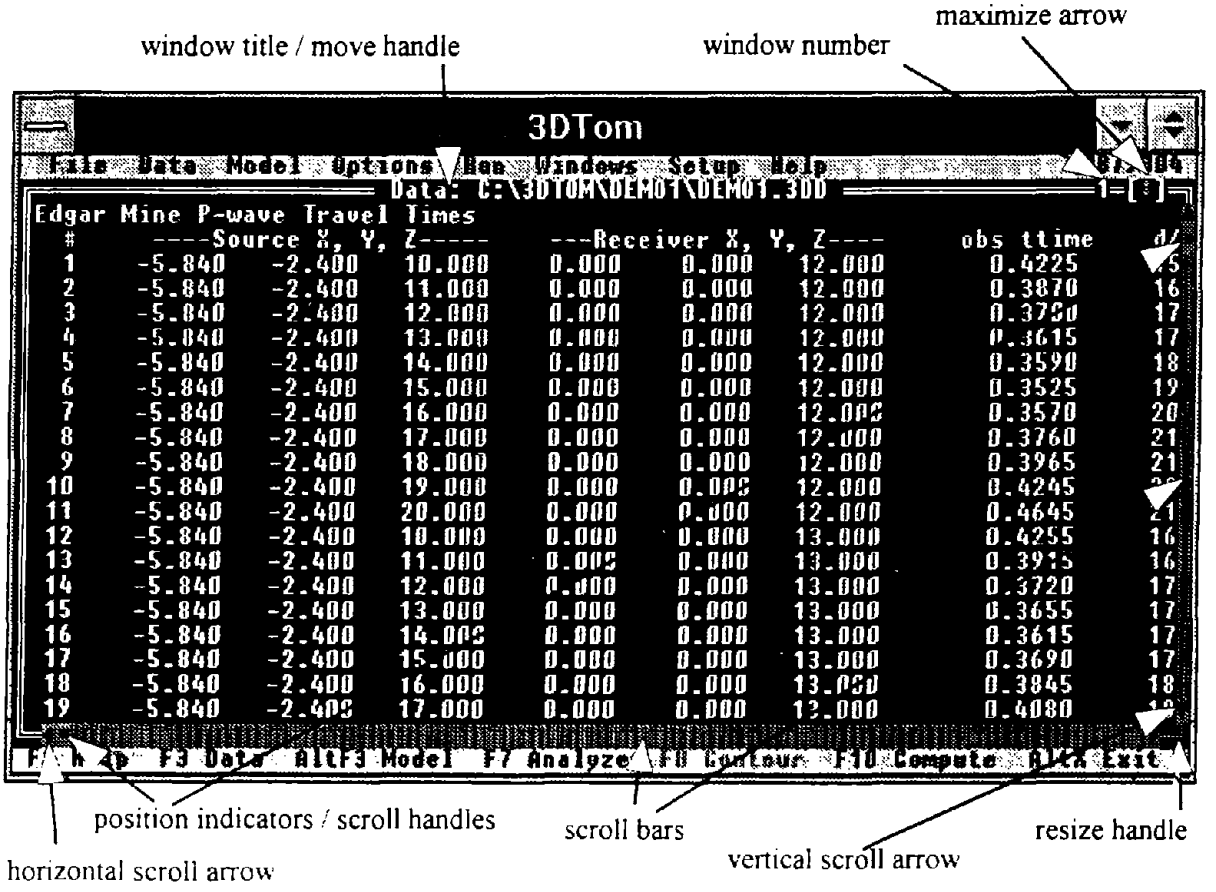

Figure 2. A text window (in this example, showing the DEMO1 data file). Mouse-sensitive visual controls are indicated.

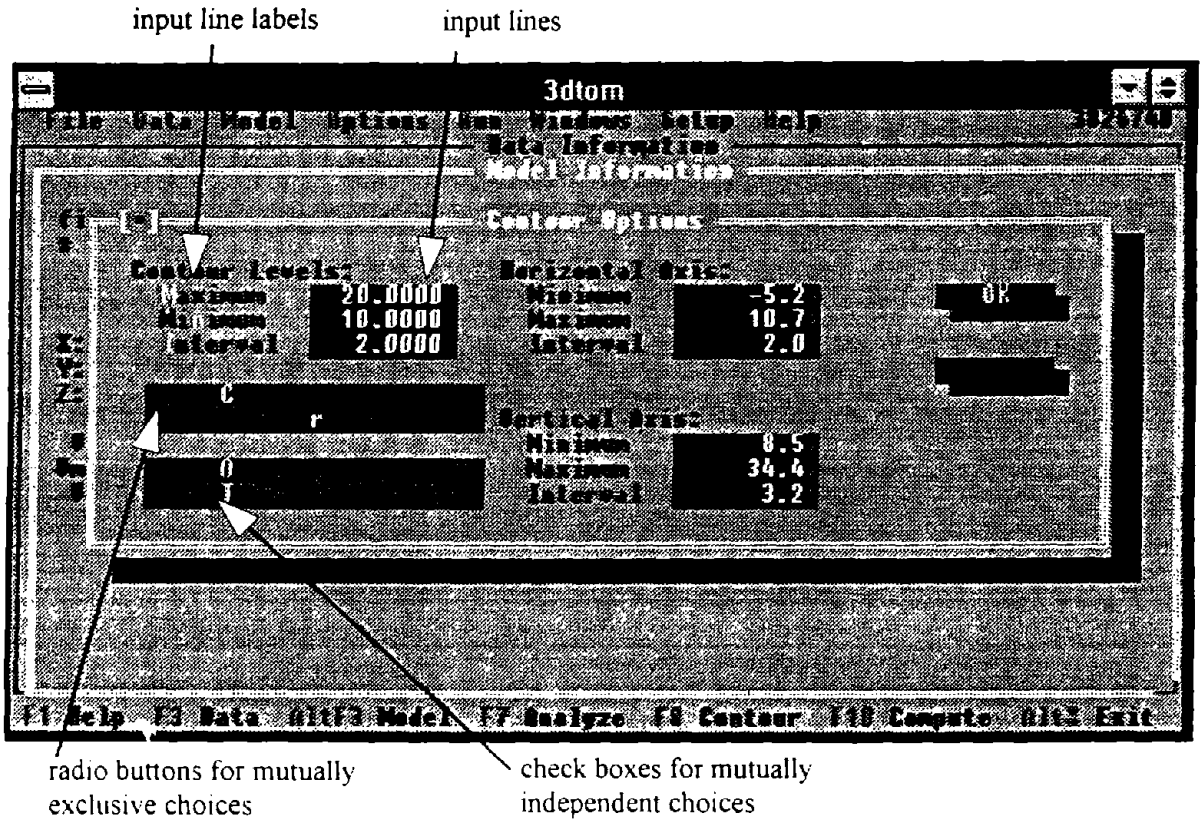

Figure 3. An example of a dialog box with several different control elements. The controls can be activated by a mouse or through the keyboard.

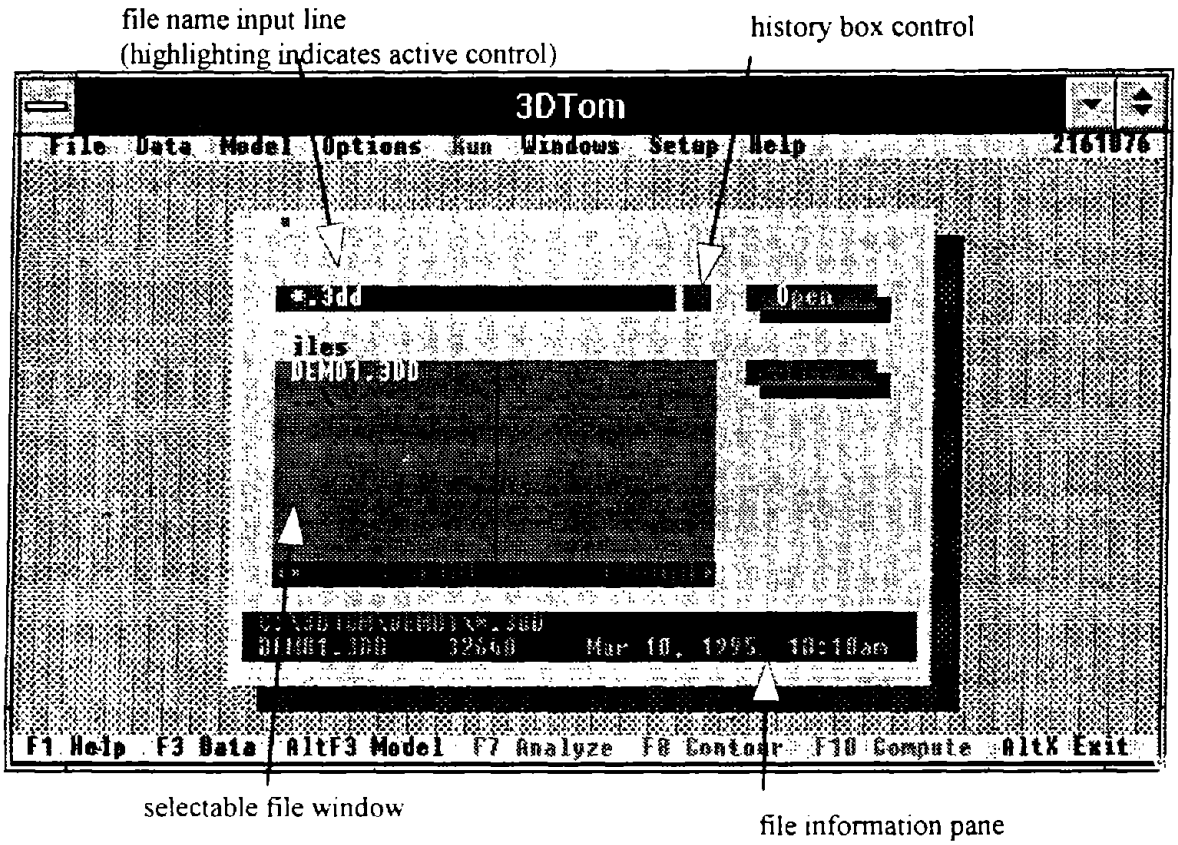

Figure 4. A file dialog box (in this example, for opening a data file).

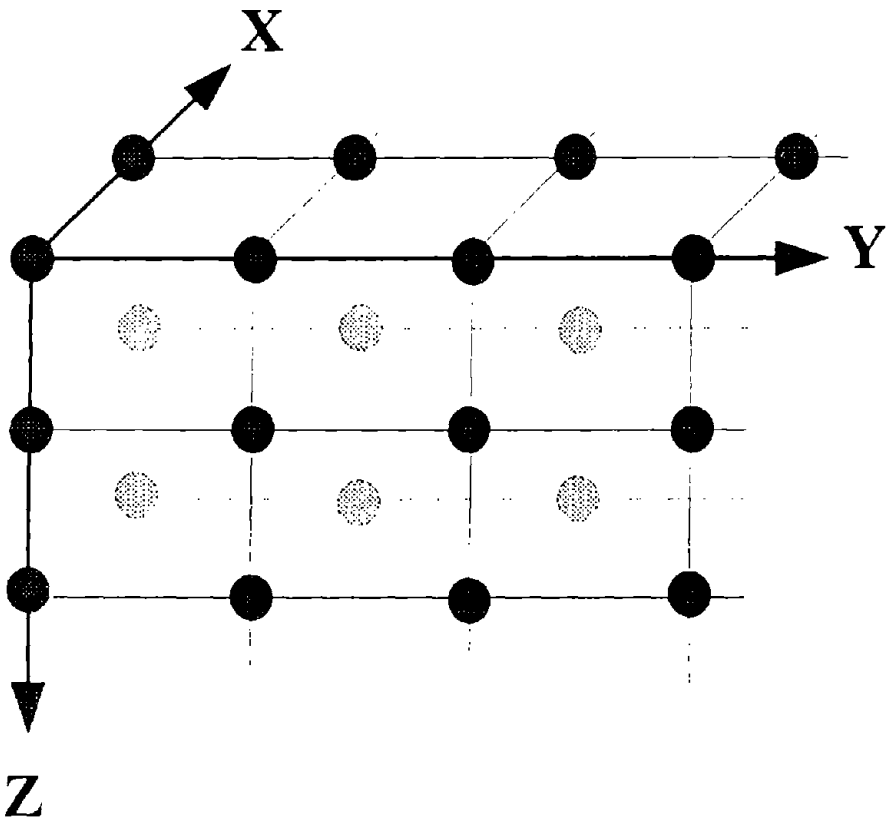

Figure 5. Schematic illustration of a portion of a 3-D grid. The grid consists of a rectilinear array of nodes (shown by circles) and intervening cells or voxels.

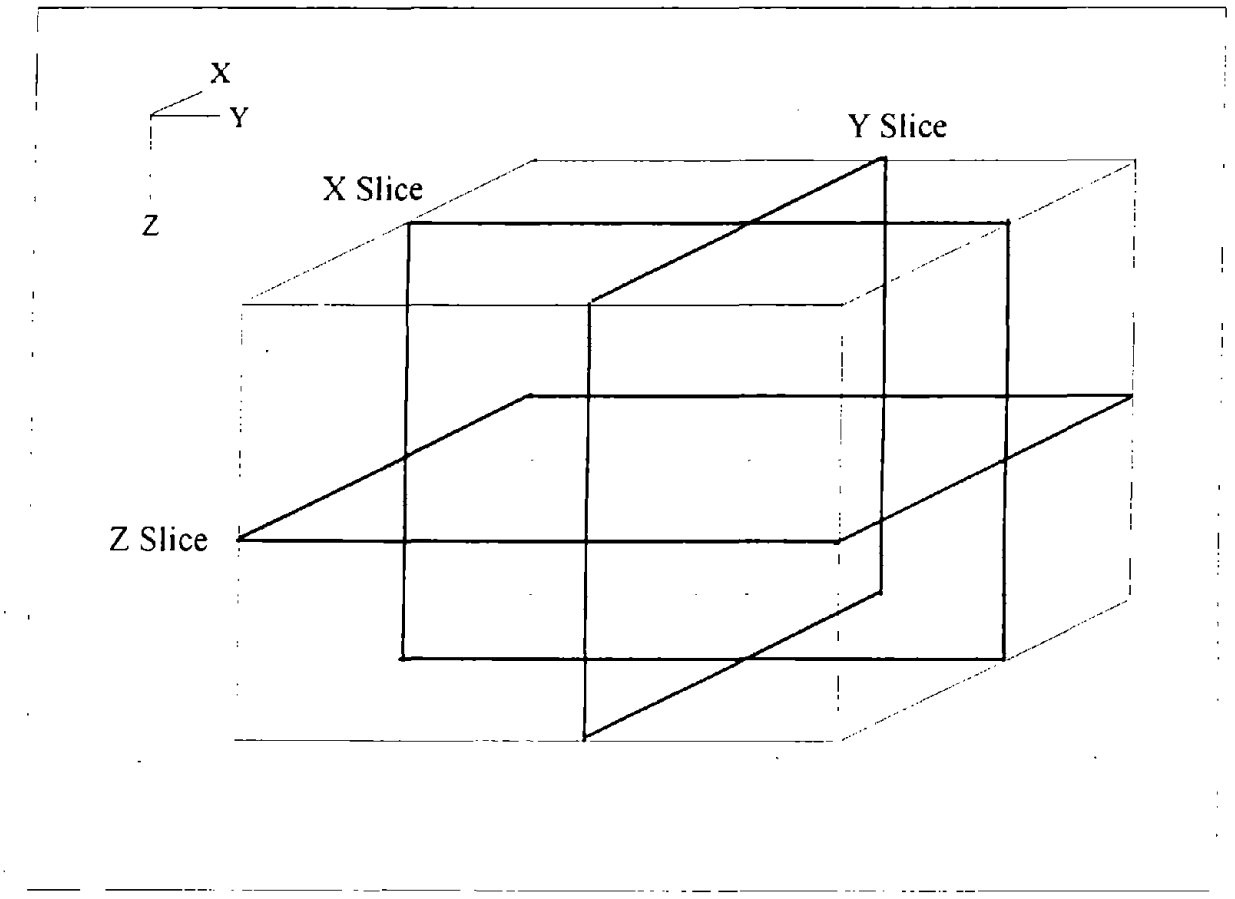

Figure 6. A model can be examined by viewing contour plots on slices perpendicular to the coordinate axes.

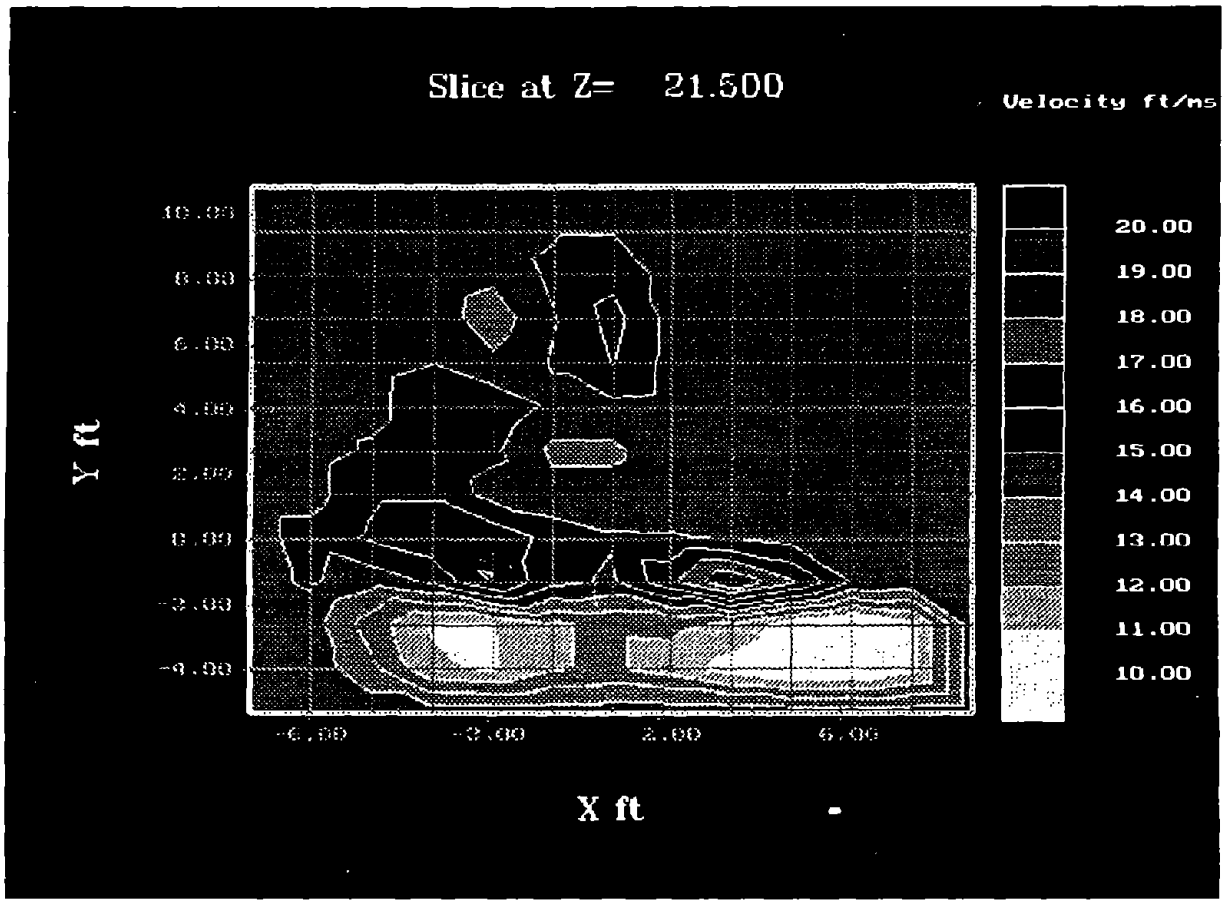

Figure 7. Example of a contour plot on a slice perpendicular to the Z axis (based on the DEMO1 data set). Grid lines are also shown in this example.

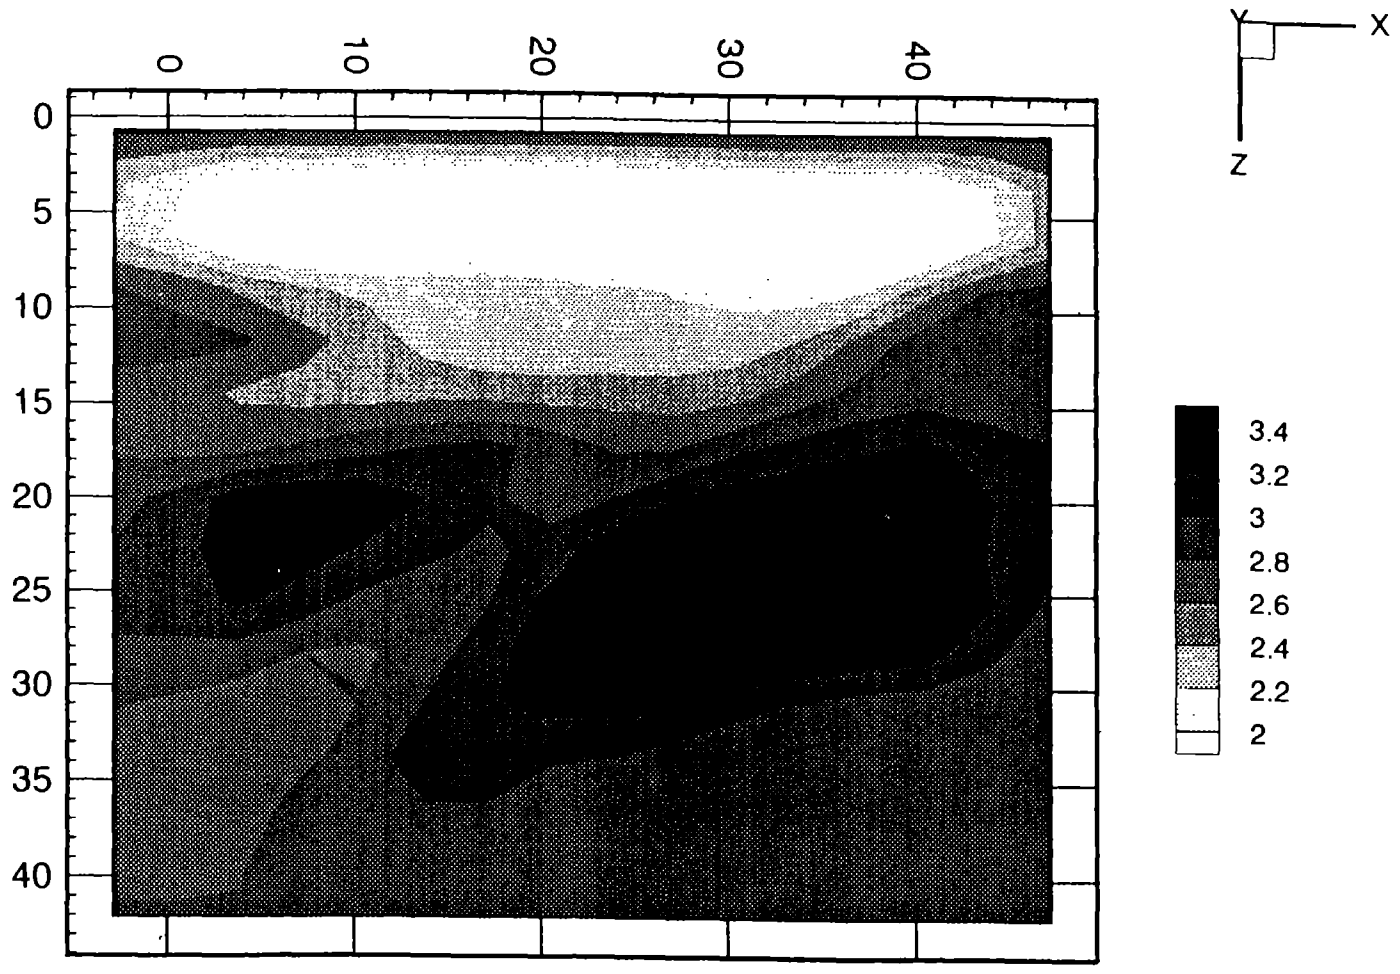

Figure 8. *Example of the effects of the simple smoothing operators.*
A, unsmoothed velocity reconstruction based on crosshole seismic data from a horizontally-layered rock mass

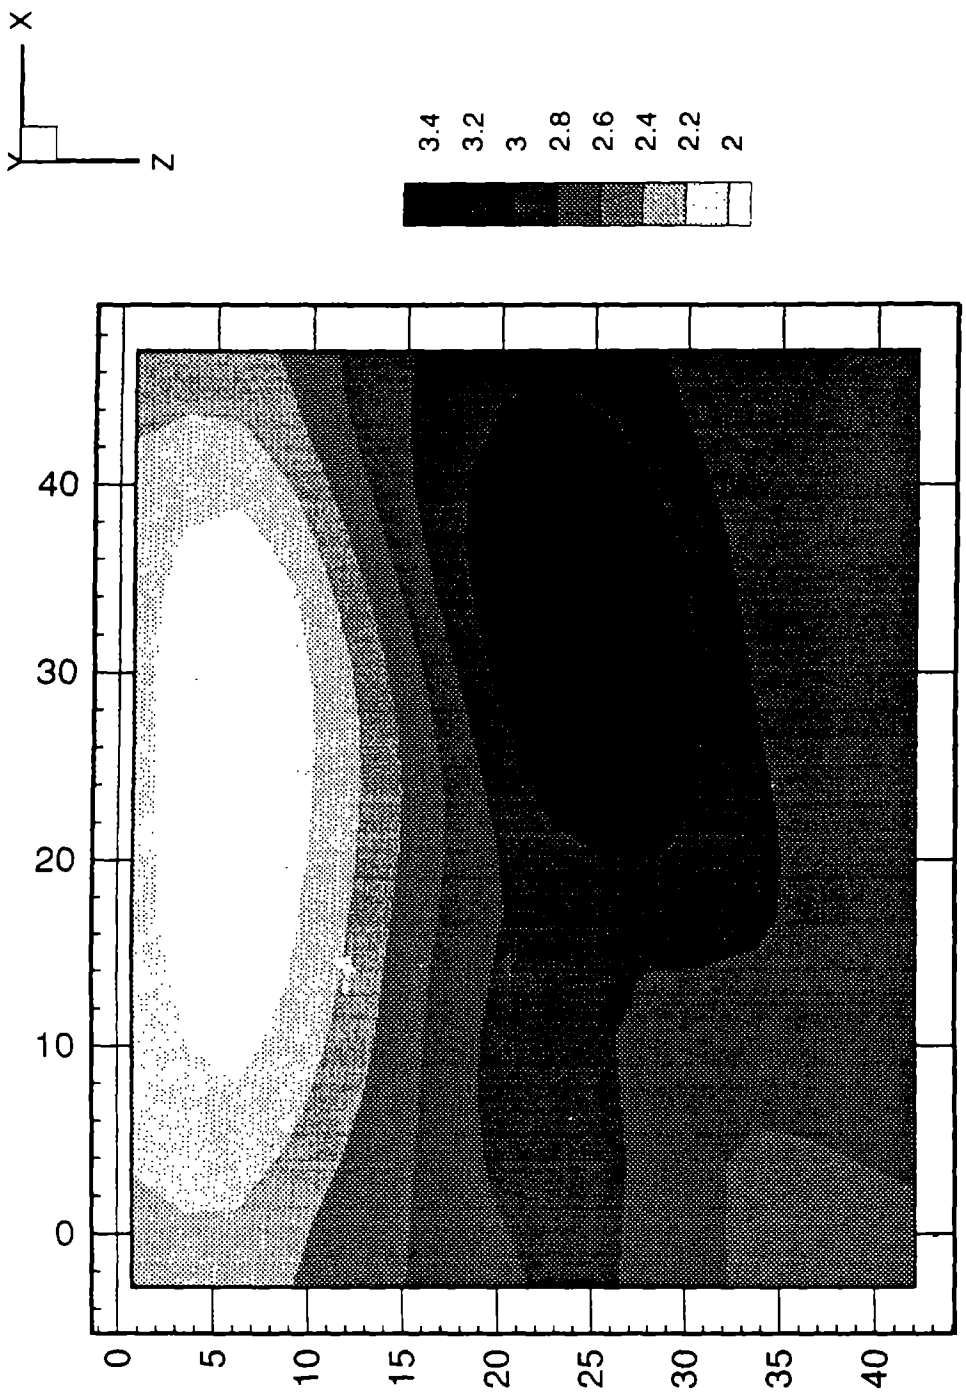

Figure 8, continued.
B, velocity reconstruction after two applications of omnidirectional smoothing.

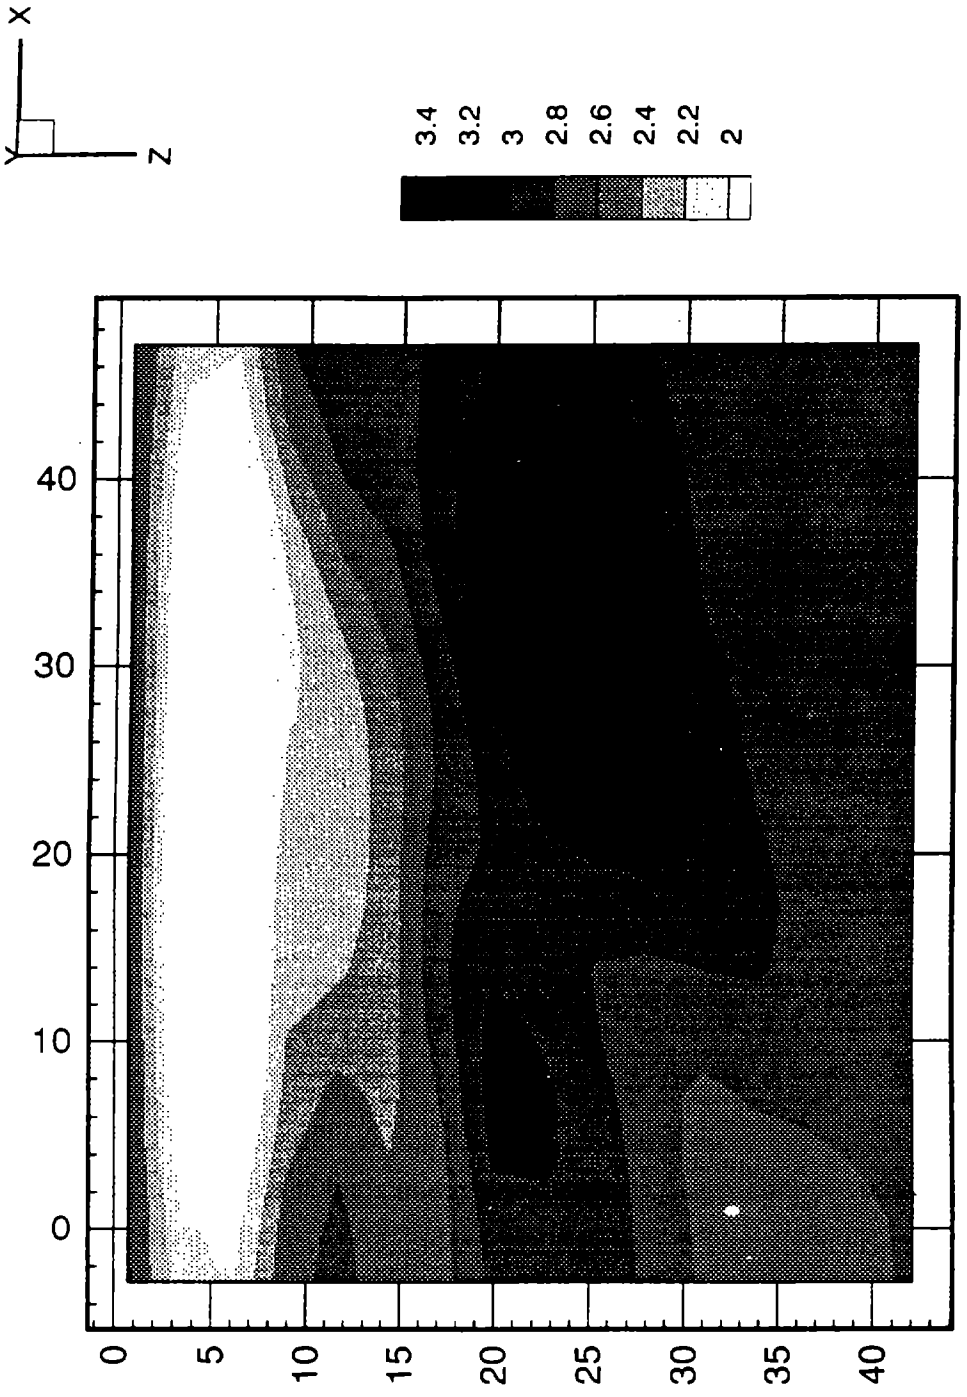

Figure 8, continued.
C, velocity reconstruction after two applications of horizontal smoothing.

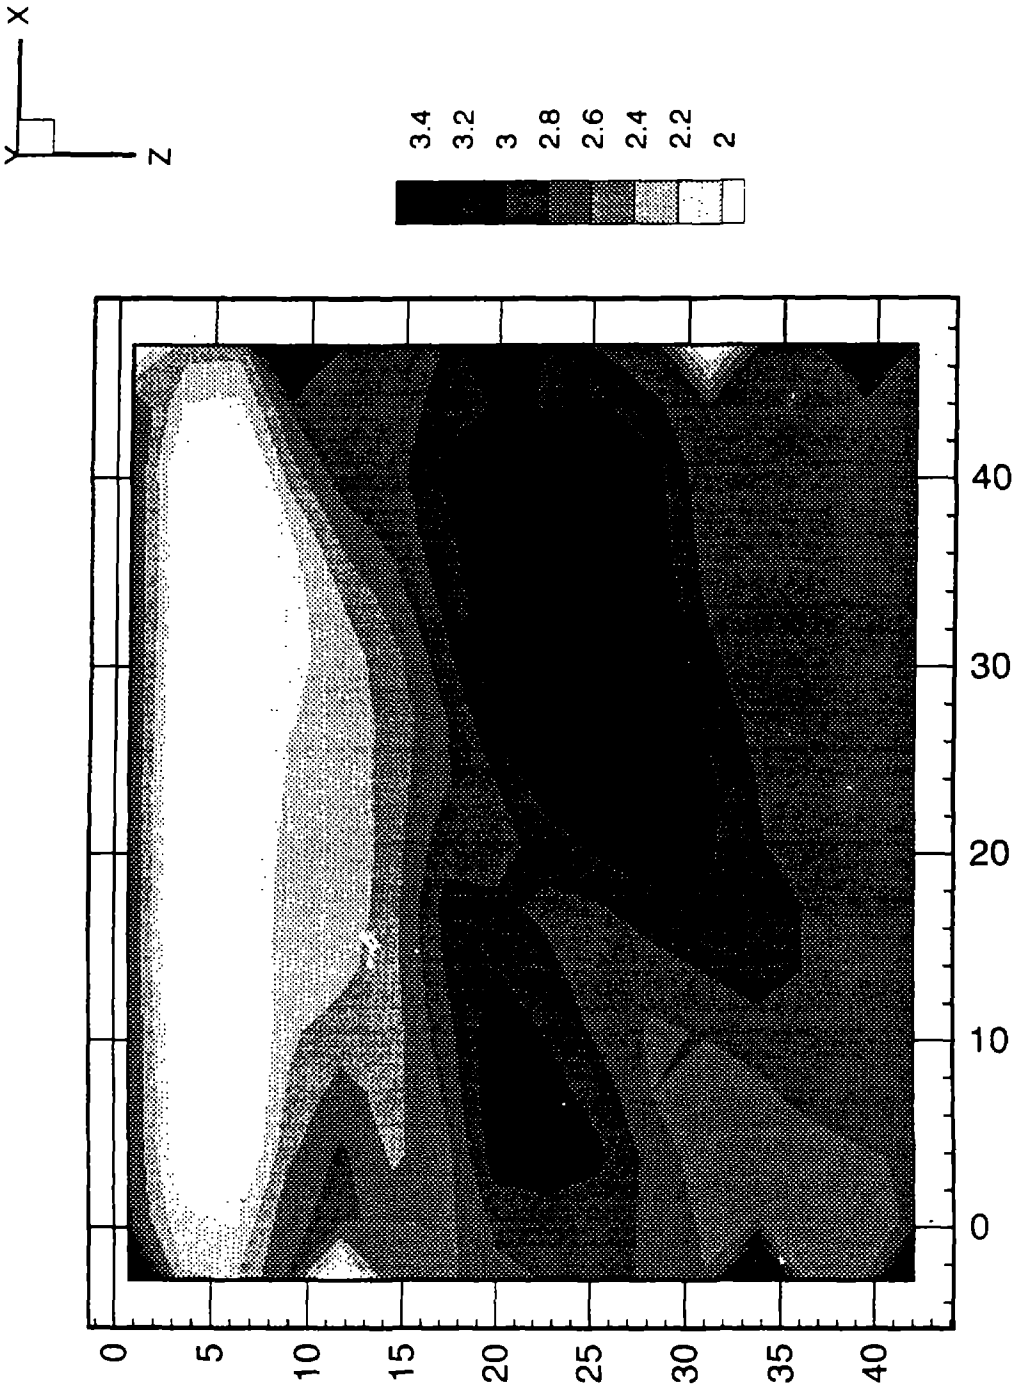

Figure 9. Example of spatial coherence filtering.
A, model of figure 8A with localized noise added along the edges. Such localized noise often results from source/receiver statics.

Figure 9, continued.
B, SC filtering removes the localized noise without altering the larger-scale velocity structure.

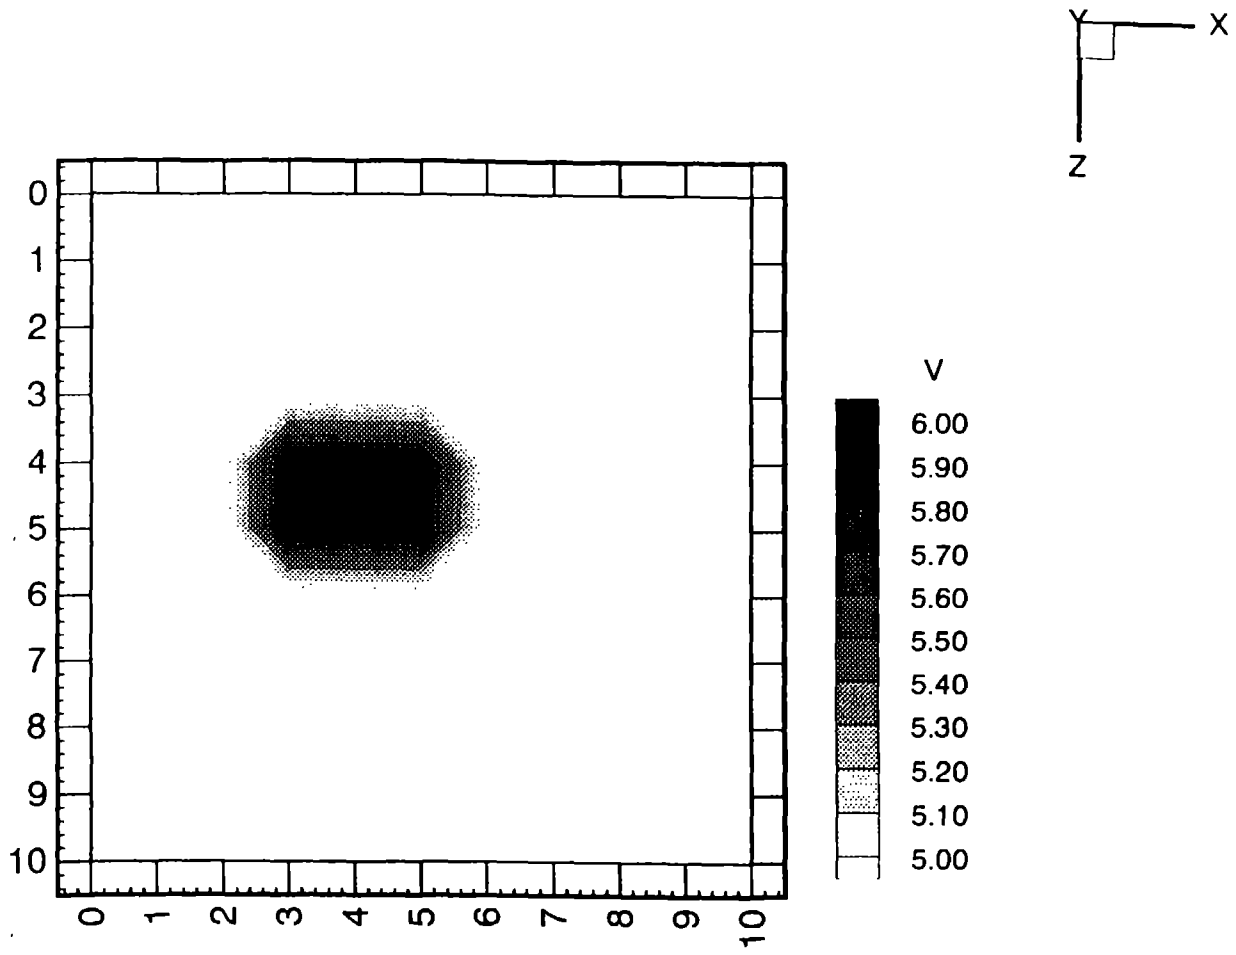

Figure 10. RAYPT example.
A, target model with a high-velocity anomaly in a uniform background. Synthetic travel times were computed for this model, with sources and receivers located along the edges of the model.

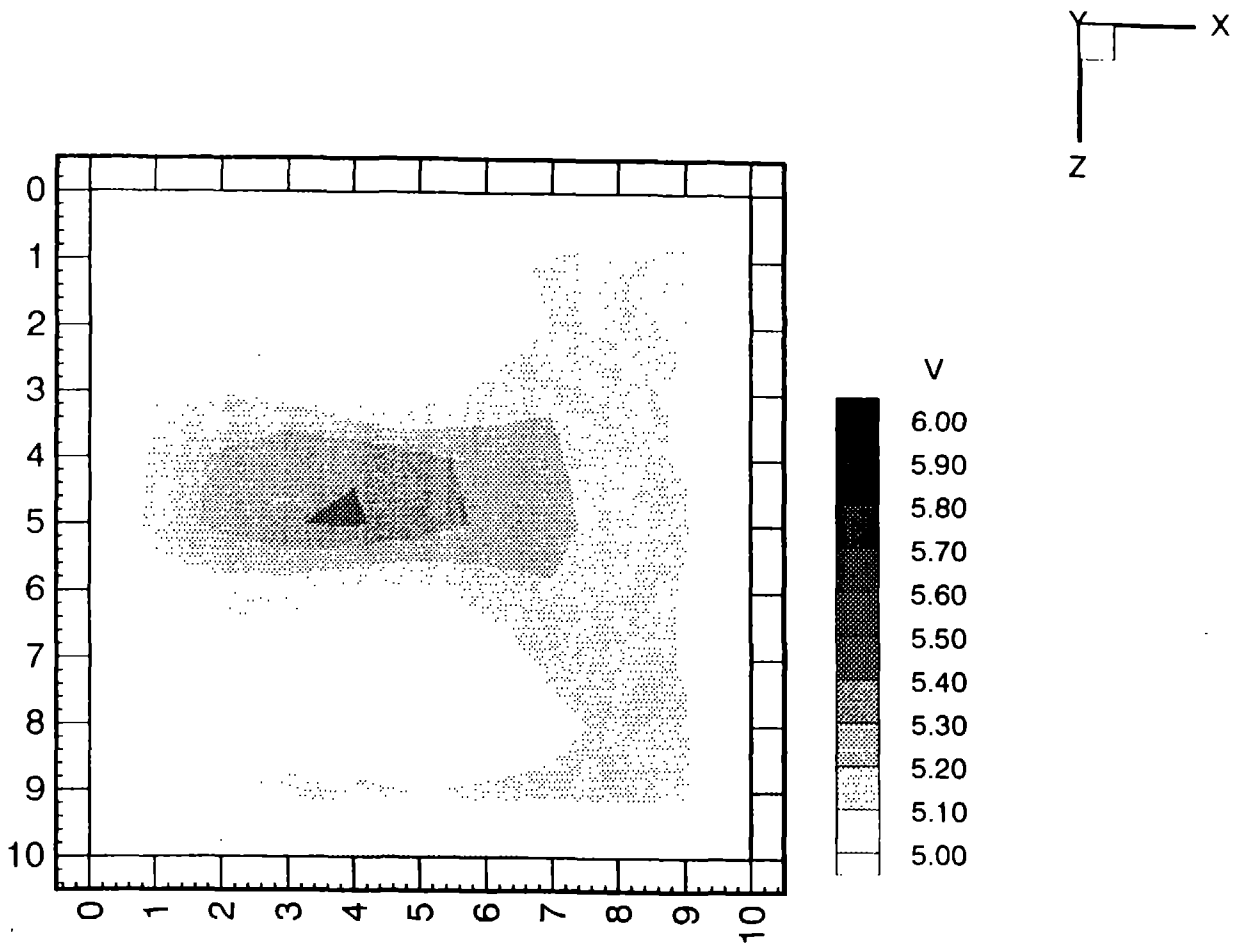

Figure 10, continued.
B, SIRT reconstruction based on the synthetic travel times shows the anomaly in a heterogeneous background.

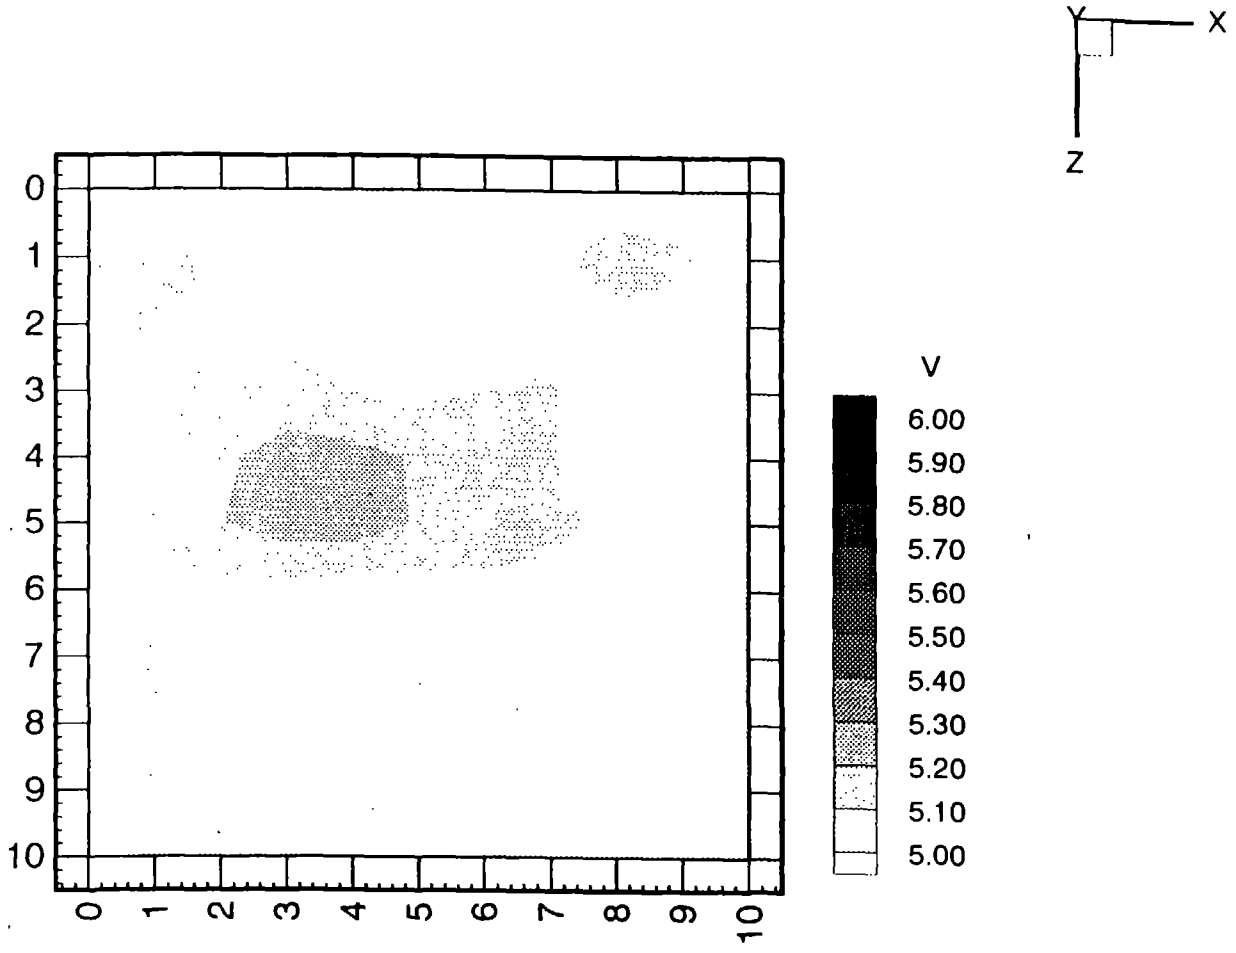

Figure 10, continued.
C, RAYPT reconstruction based on the synthetic travel times shows the anomaly in a relatively uniform background.

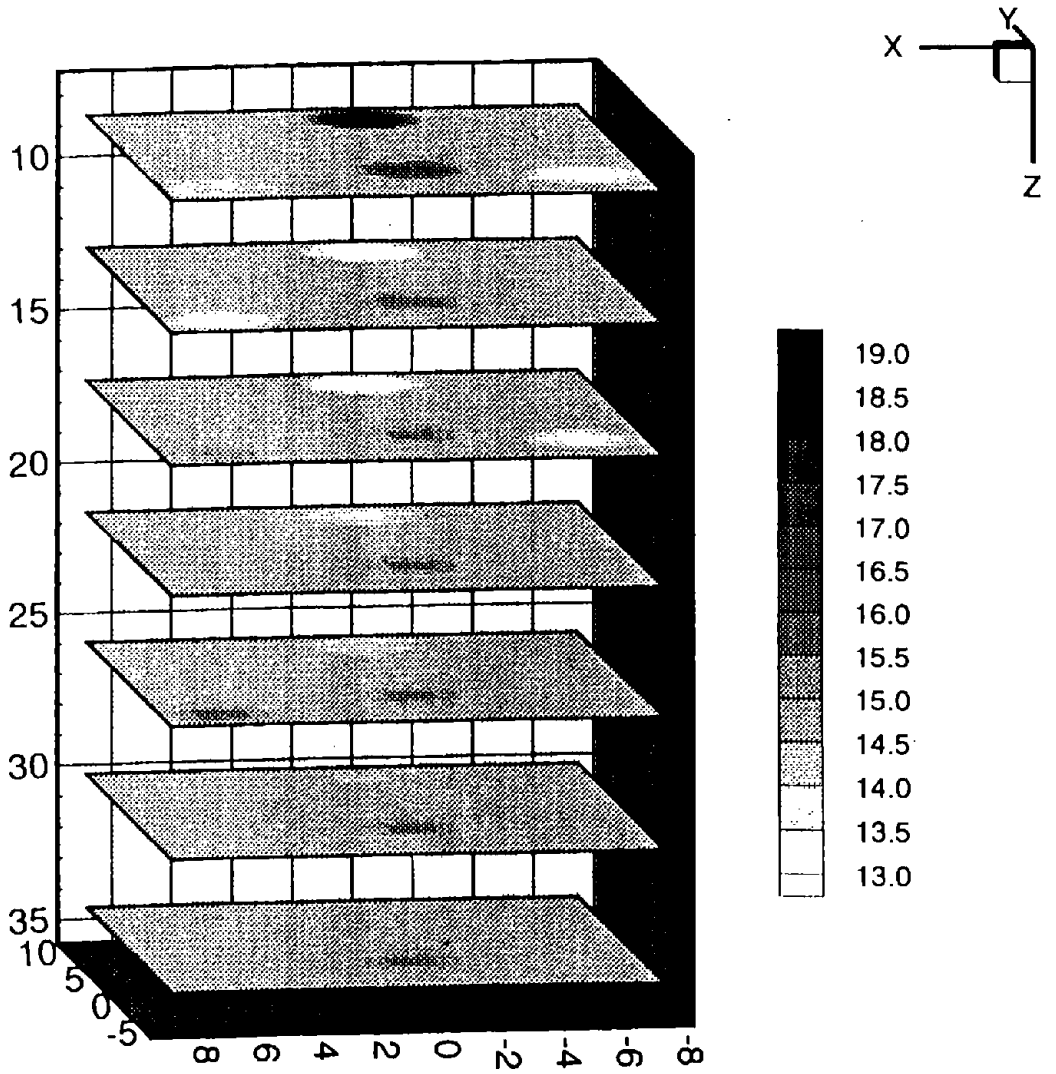

Figure 11. Application of borehole velocity logs, using hypothetical P-wave velocities for the boreholes in the DEMO1 data set. A, borehole information applied to nearest neighbor nodes.

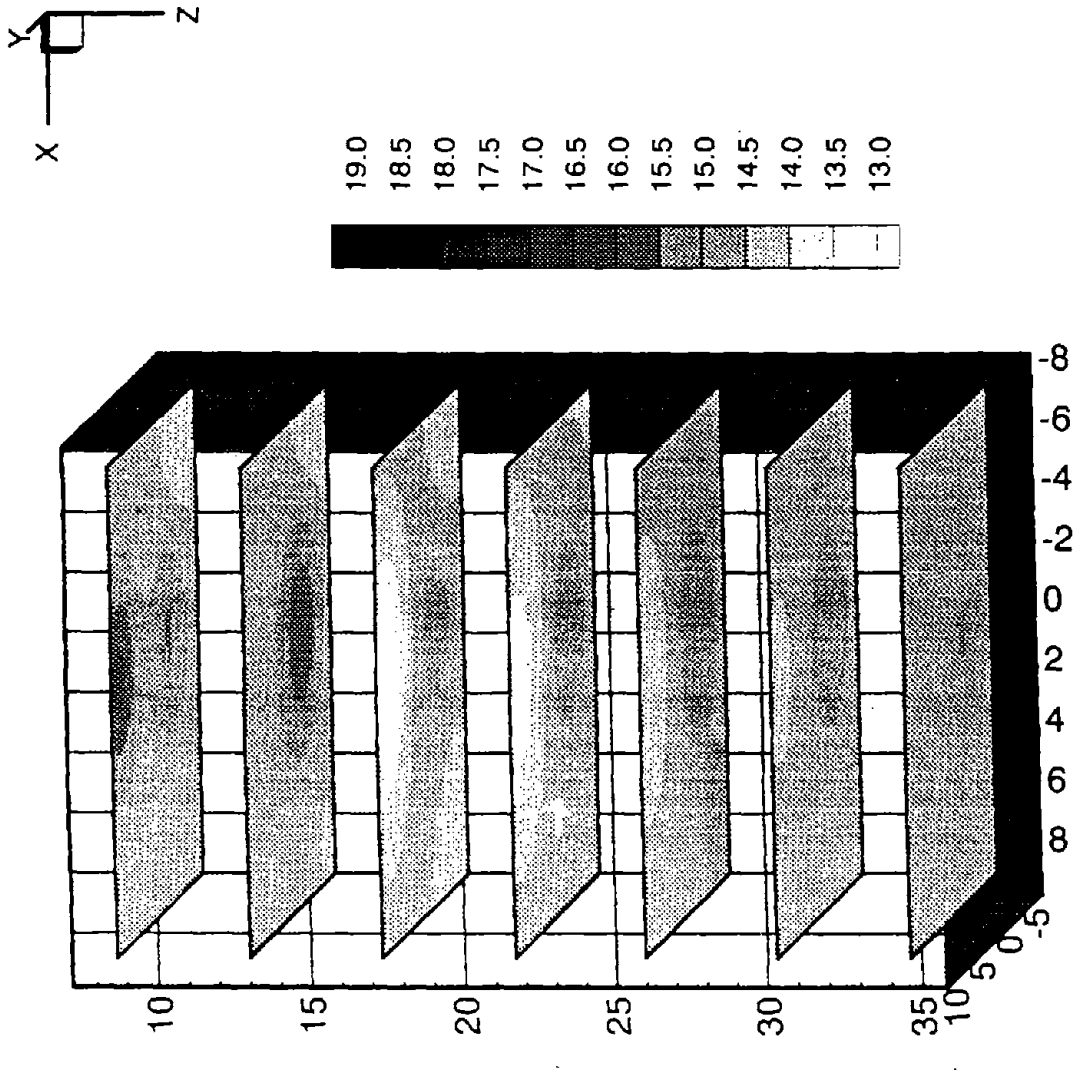

Figure 11, continued.
B, borehole information applied to all nodes with inverse distance weighting.

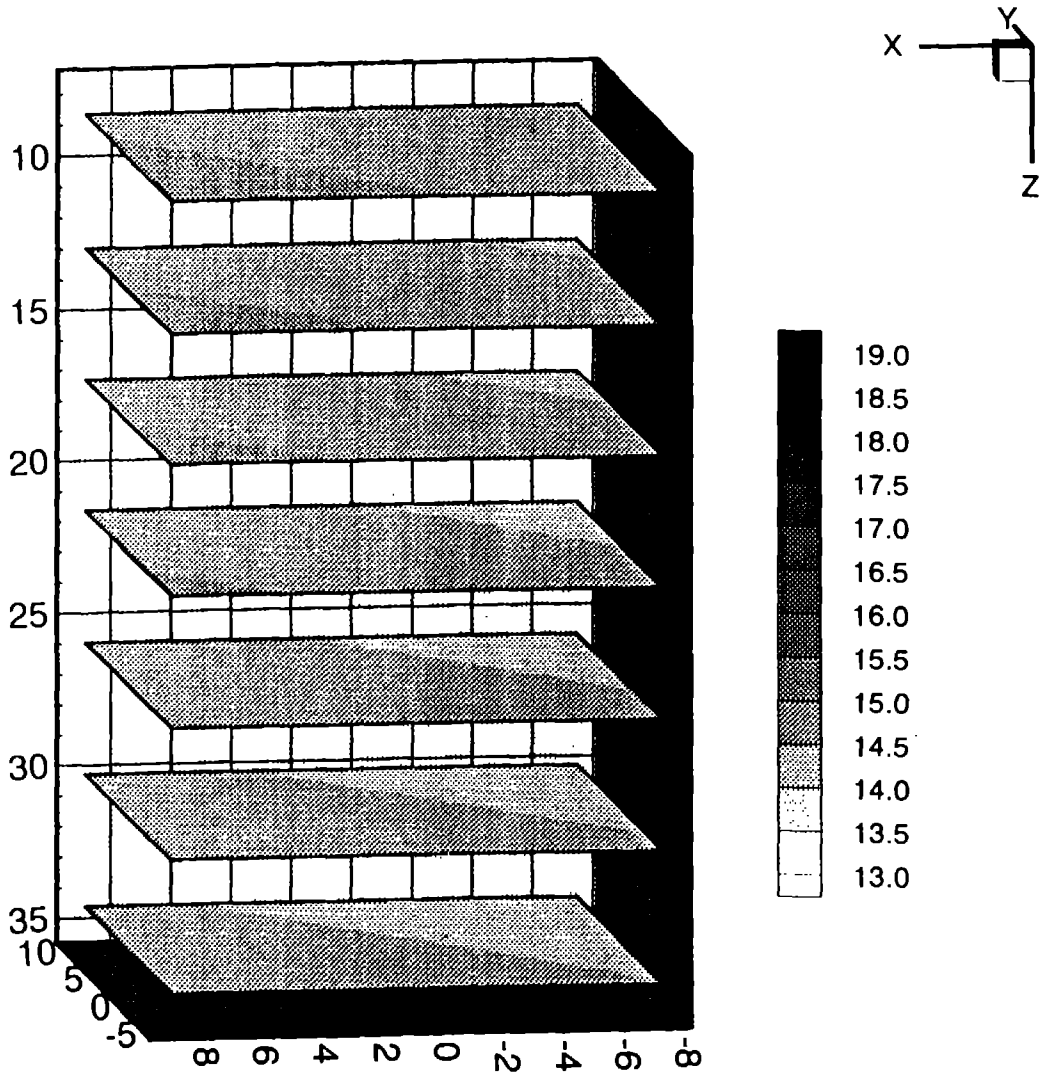

Figure 11, continued.
C, borehole information applied to all nodes with least-squares polynomial fit of degree 1 (uniform gradient).

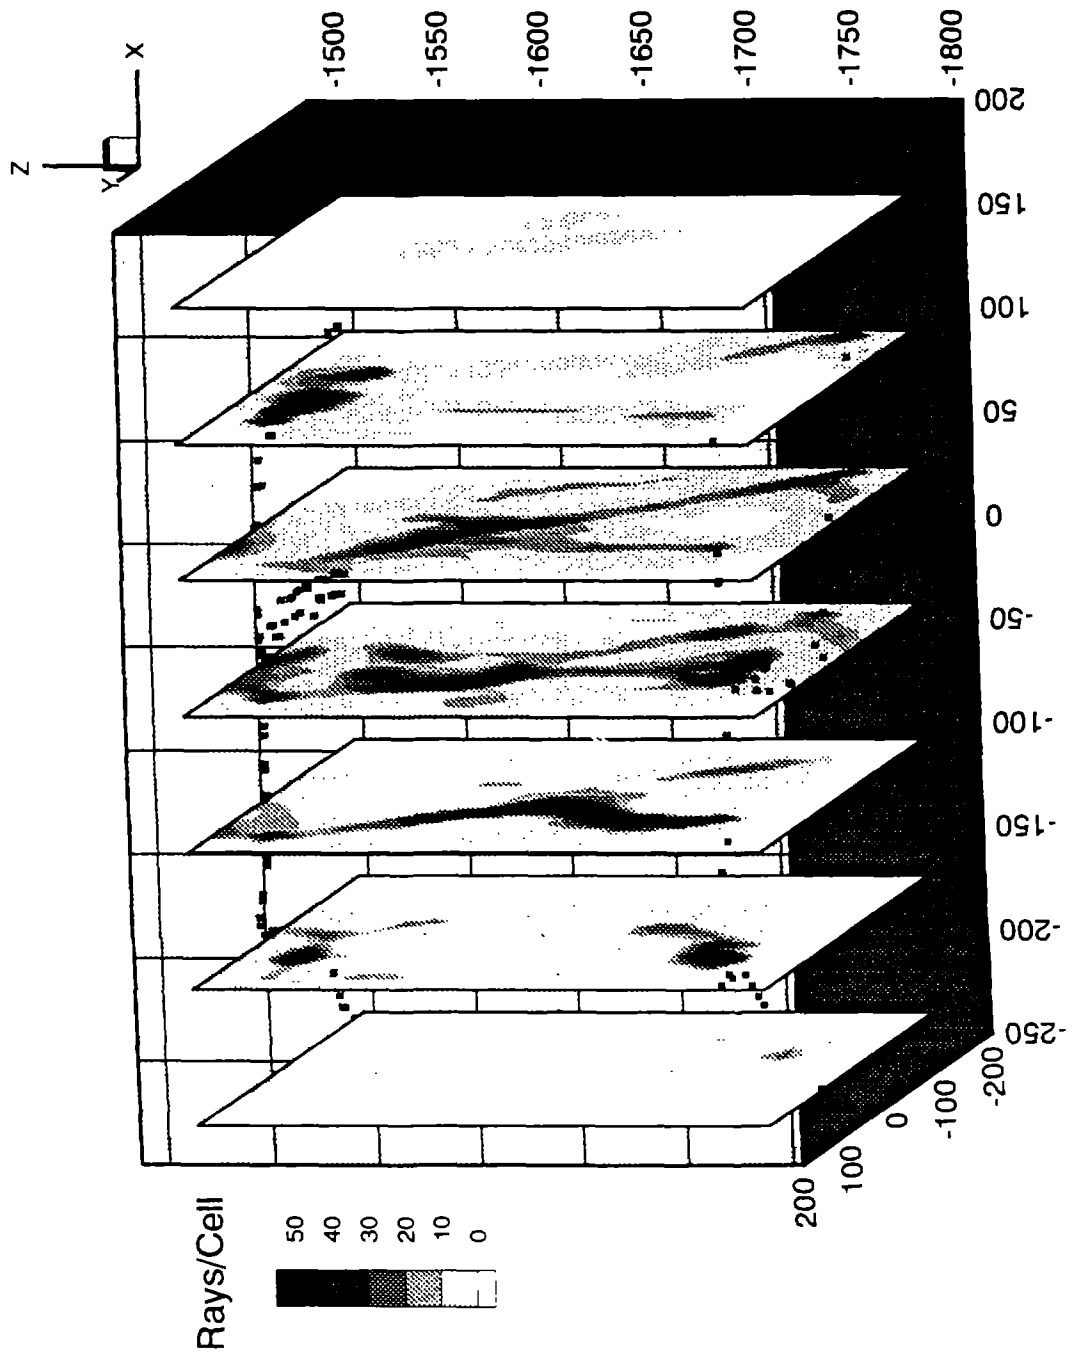

Figure 12. Ray path sampling density for the DEMO2 data set, using the default grid and curved rays. Small squares indicate source and receiver locations.

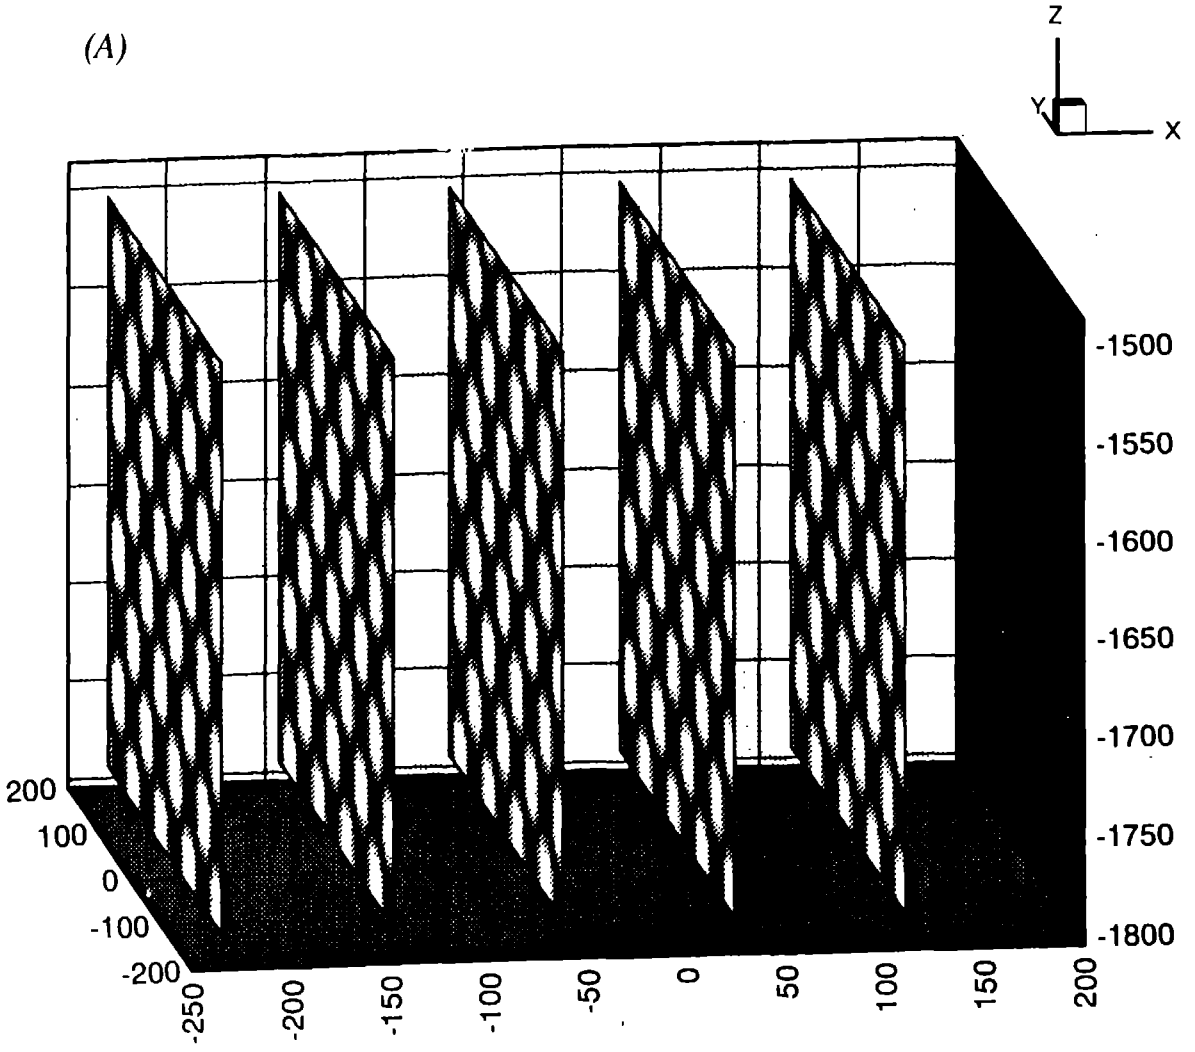

Figure 13. Checkerboard resolution test for the DEMO2 data set. A, target model. Synthetic data were computed using this model and the source and receiver locations from the real data.

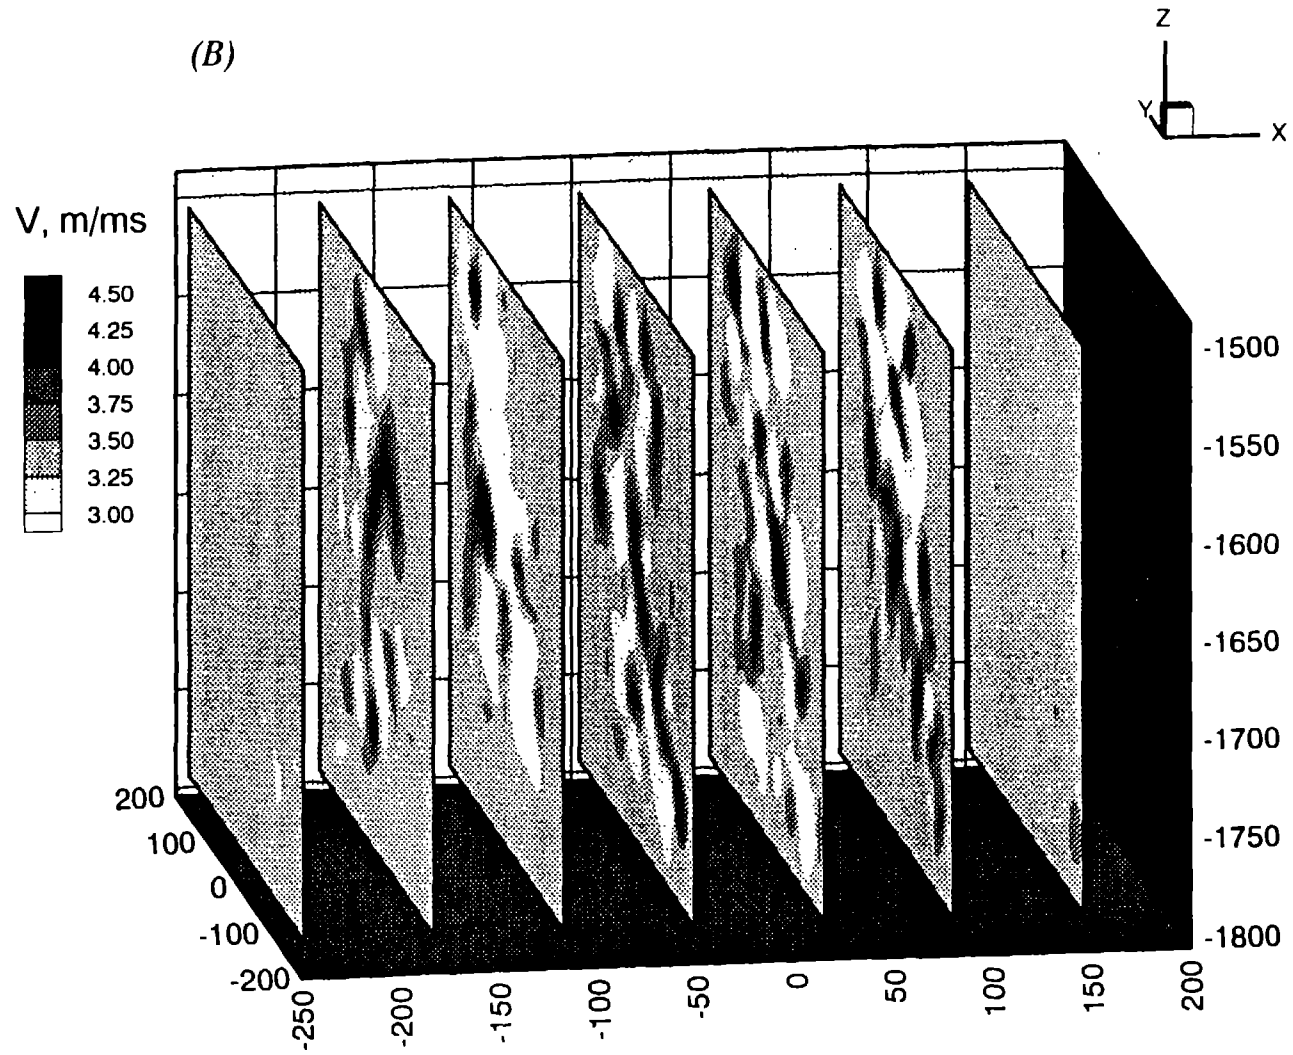

Figure 13, continued. B, reconstruction. The checkerboard variation is most visible in the regions with high sampling density - these are the regions of maximum resolution for this data set.

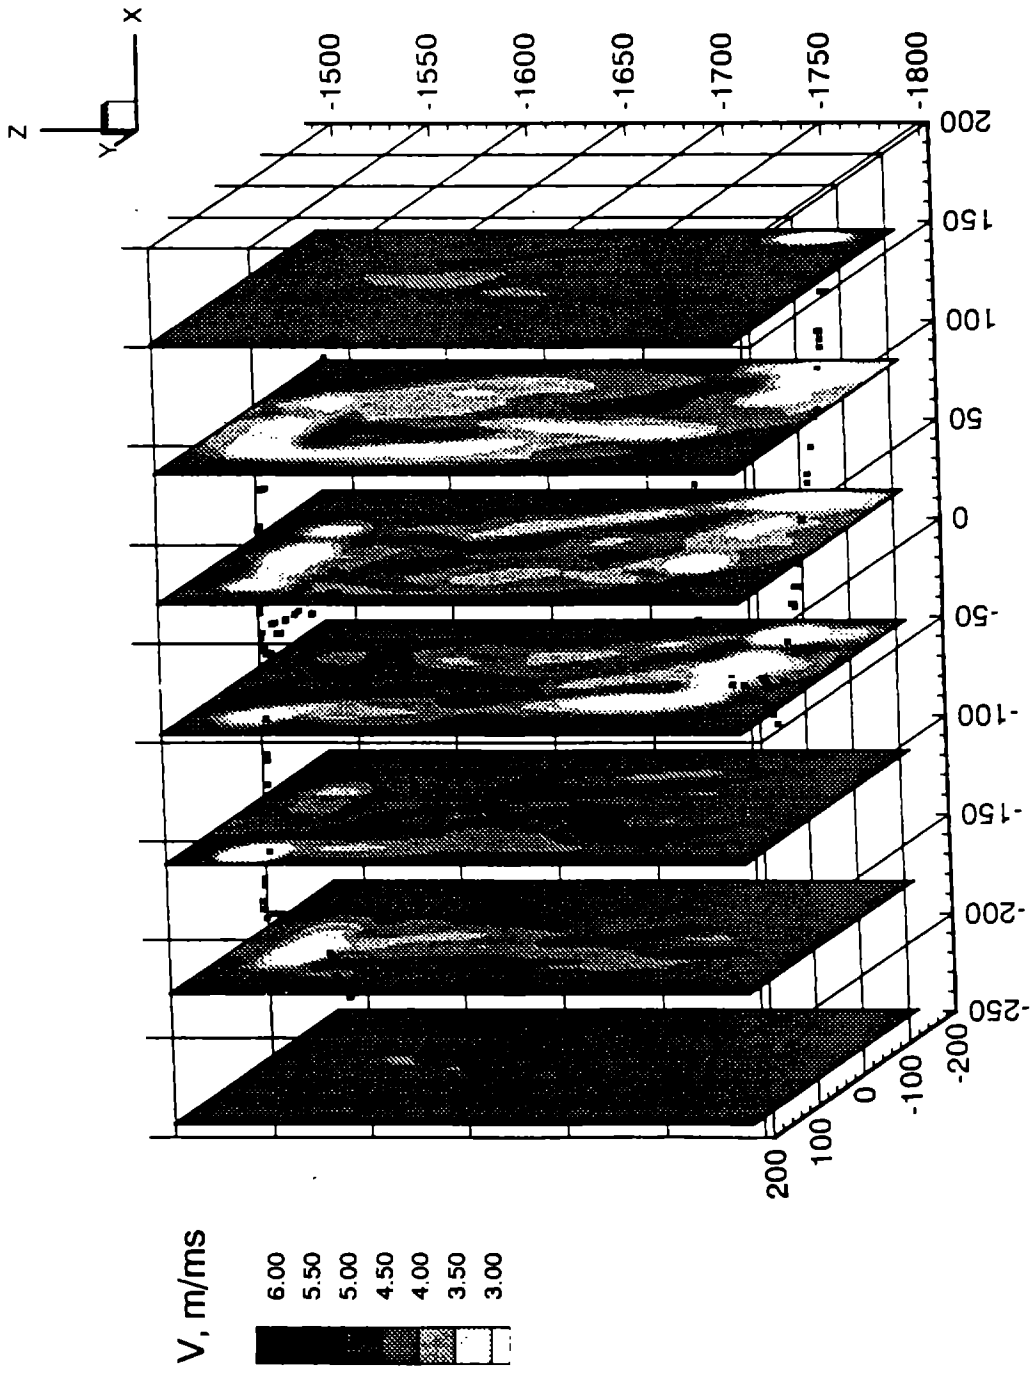

Figure 14. Average and standard deviation of ten reconstructions derived from different starting models. A, mean reconstructed velocities are generally low in and around the mine openings, and higher in the intervening rock mass.

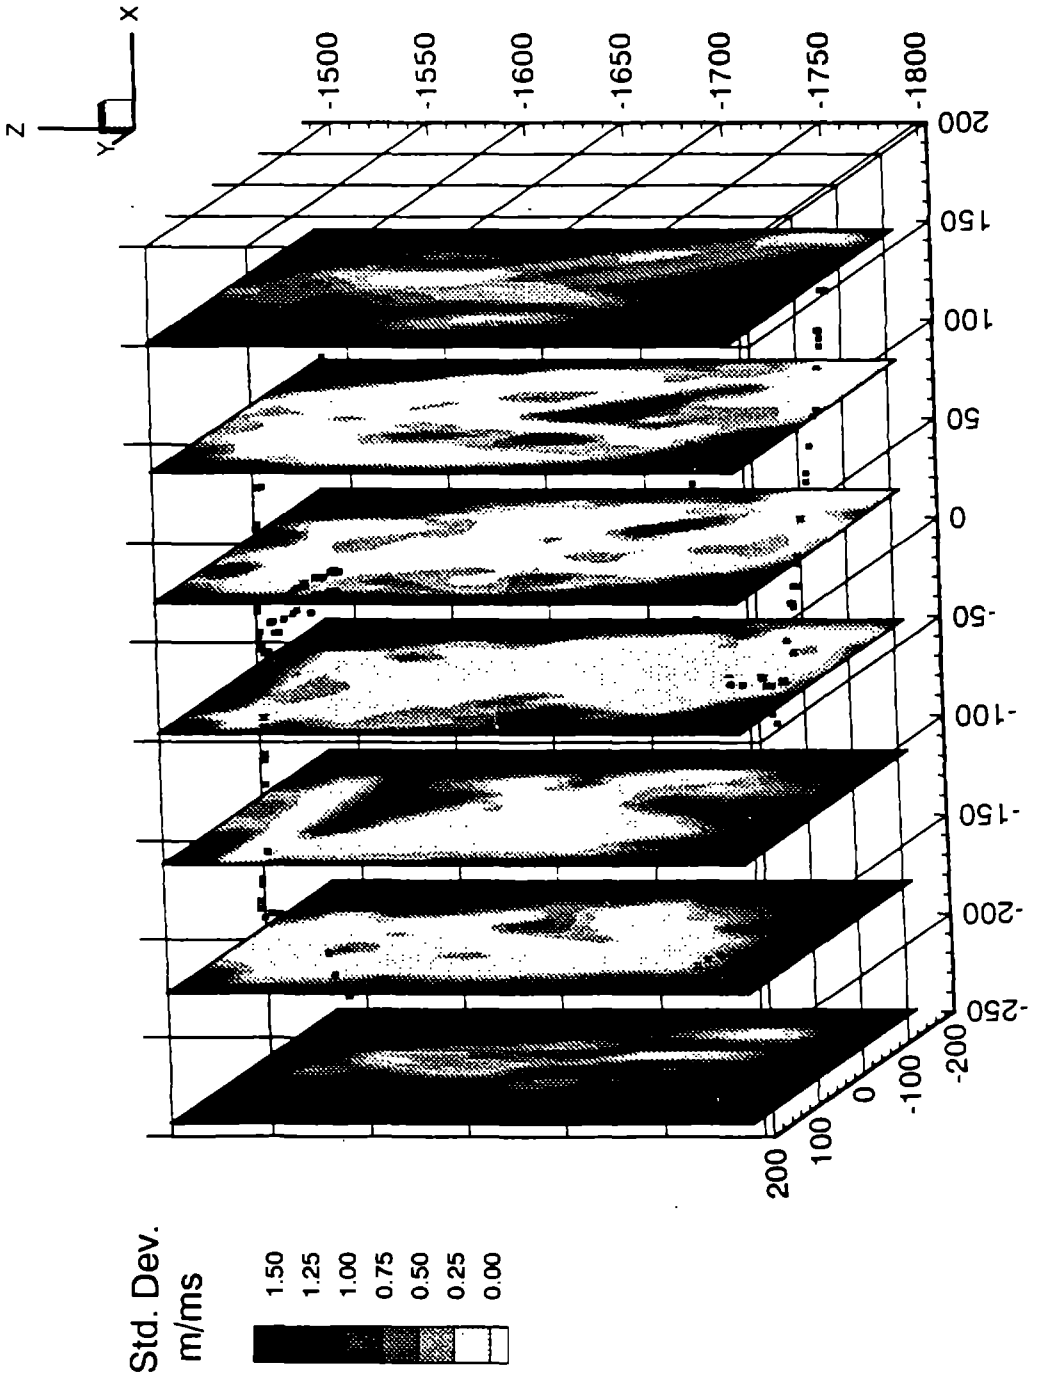

Figure 14, continued. B, low standard deviations show regions where the final model is insensitive to details of the starting model - these are the regions where the reconstruction is most robust and reliable.

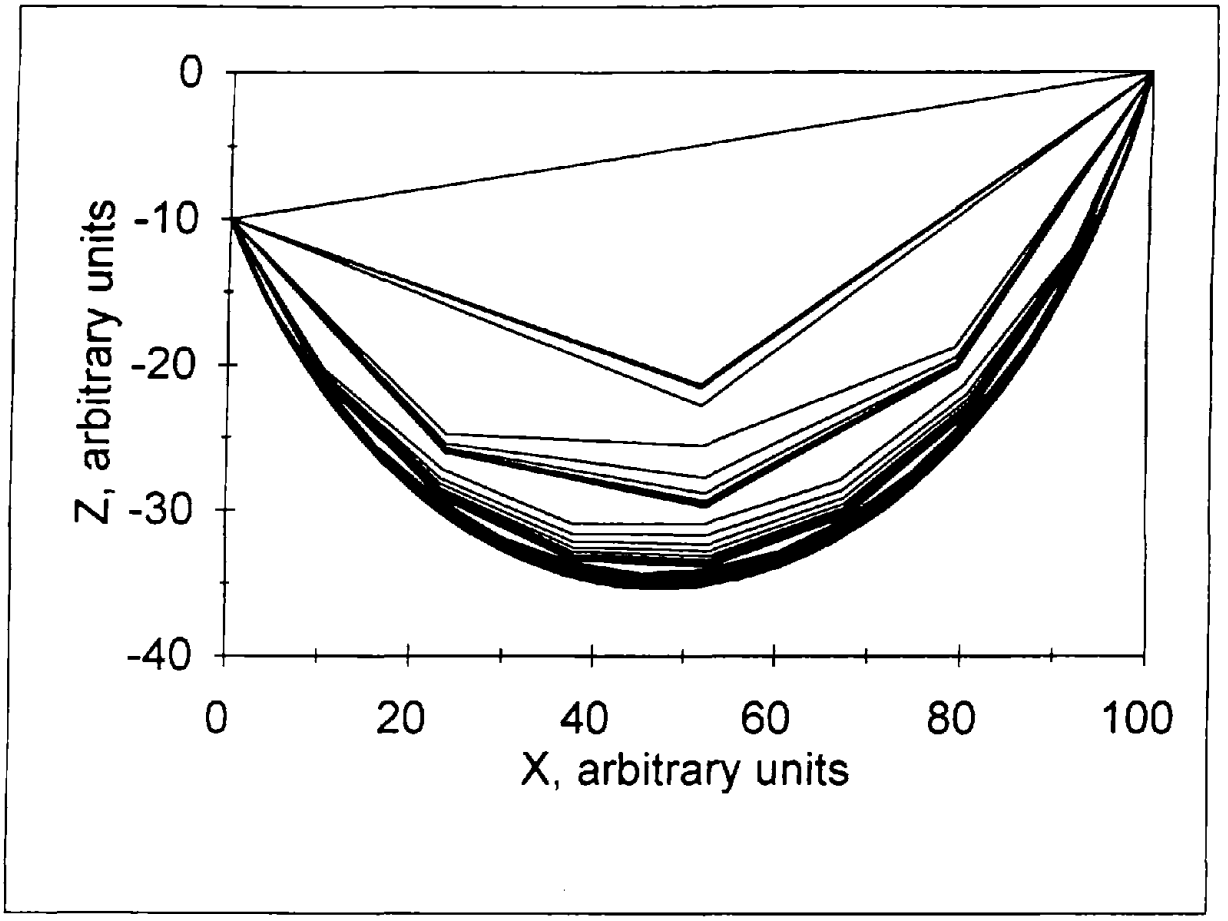

Figure 15. Two-dimensional illustration of the ray bending method, showing successive approximations to the raypath joining a source and receiver through a velocity model with a uniform vertical velocity gradient. X and Z are spatial coordinates.

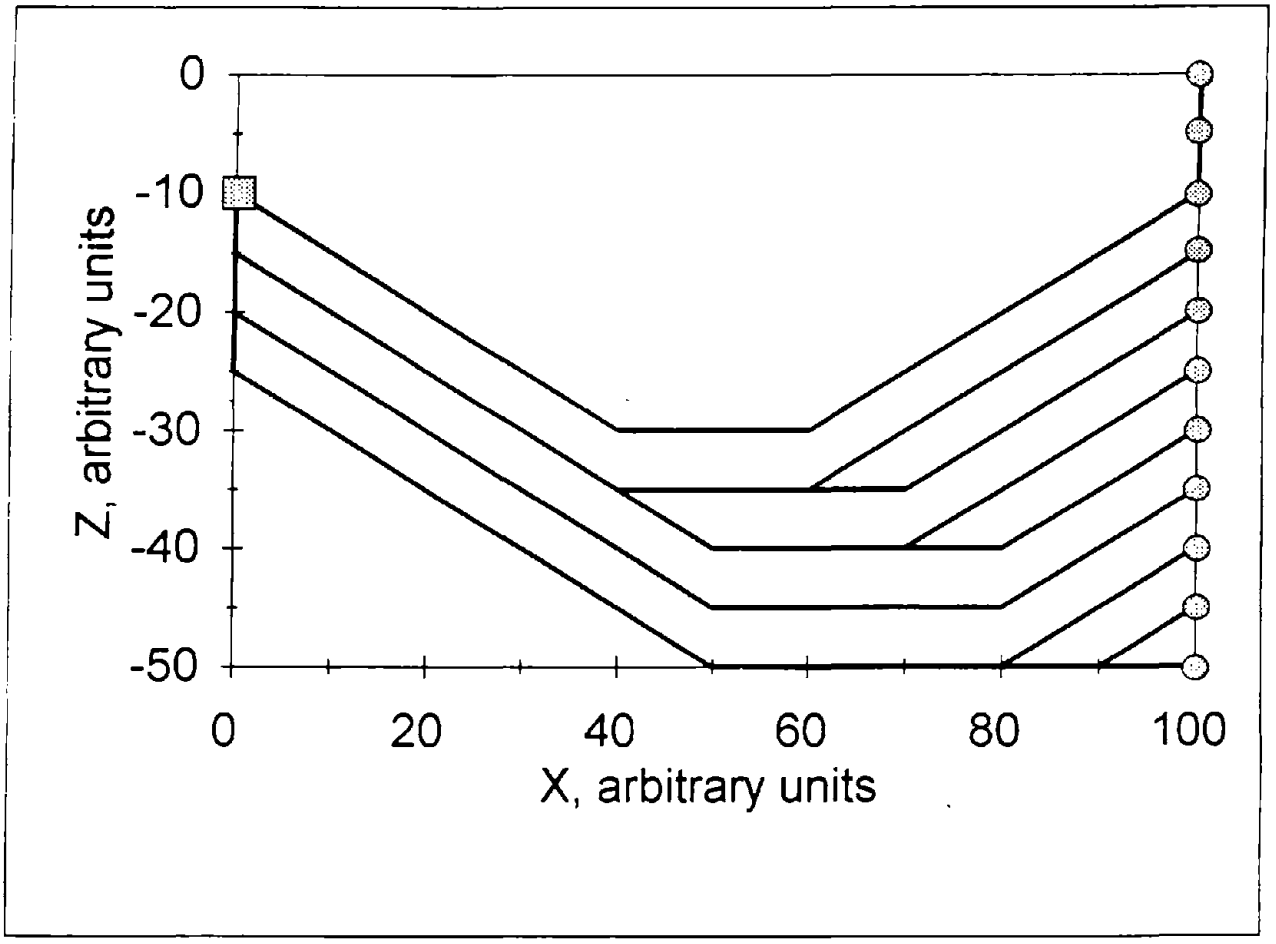

Figure 16. Rays traced using the network approach. First-arrival raypaths are shown for a source (square) on the left to multiple receivers (circles) on the right. The velocity increases linearly with depth, as in Figure 11.

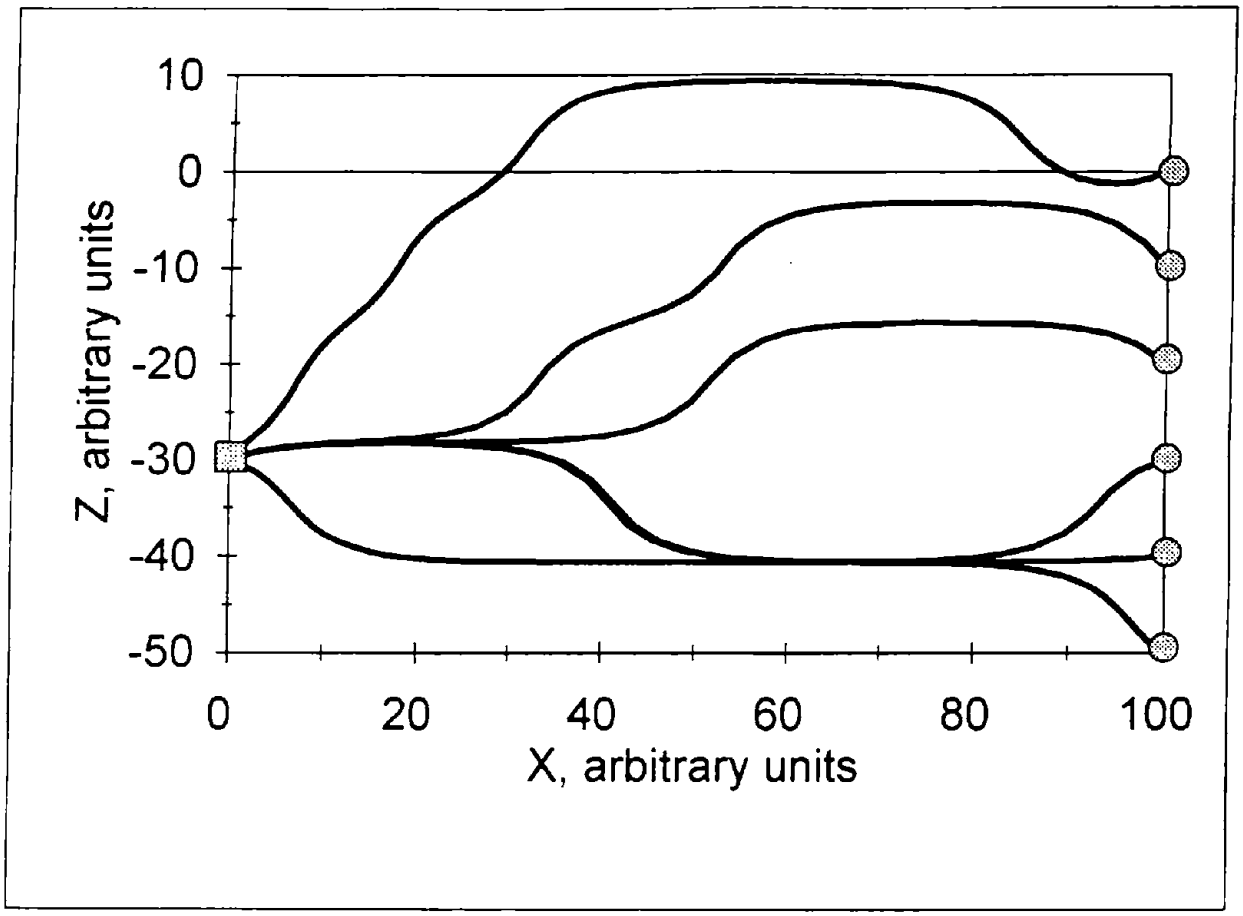

Figure 17. Rays traced using the bending approach. First-arrival raypaths are shown for a source (square) on the left to multiple receivers (circles) on the right. The velocity varies sinusoidally with depth; four high-velocity layers alternate with slow layers between 0 and 40 depth units.

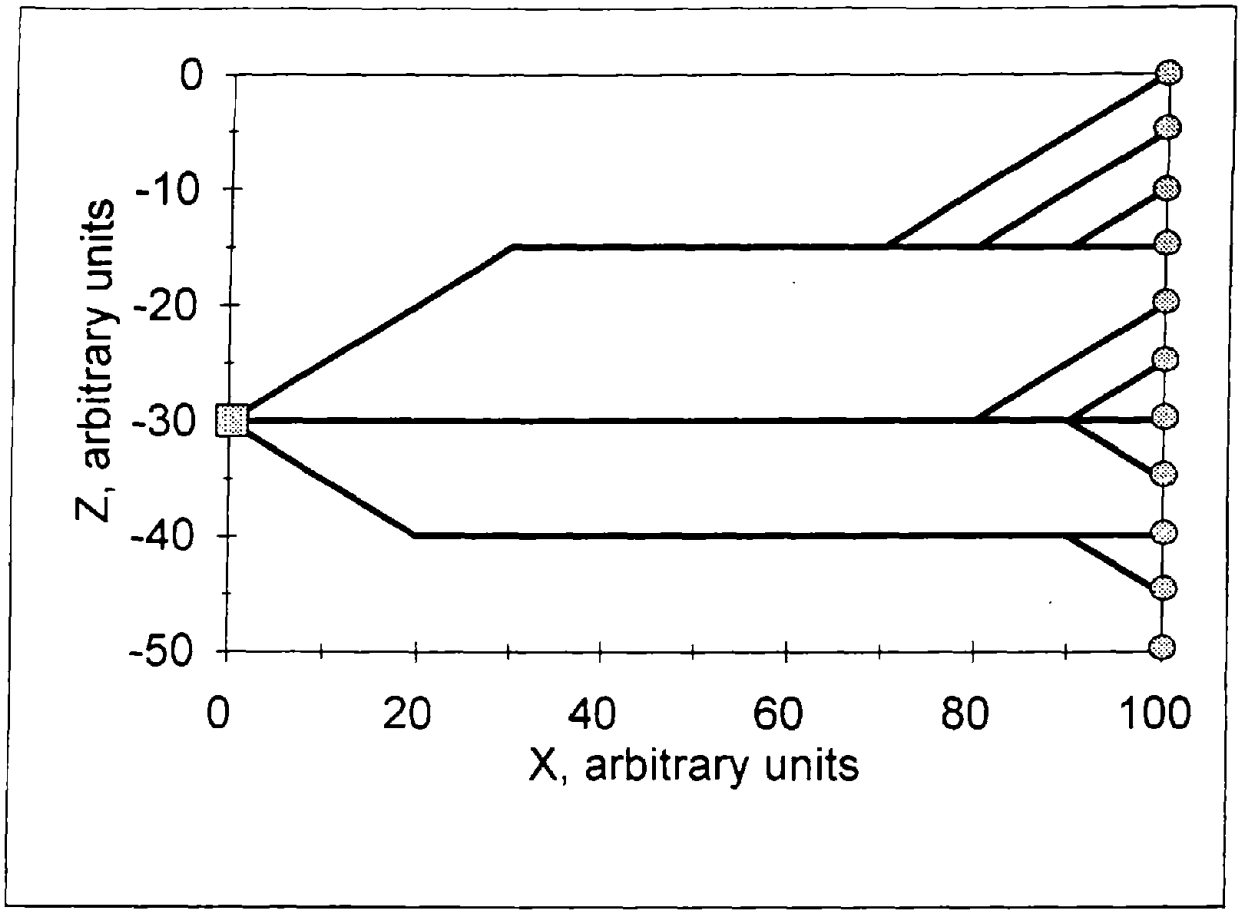

Figure 18. Rays traced using the network approach for the model with velocity varying sinusoidally with depth; four high-velocity layers alternate with slow layers between 0 and 40 depth units.

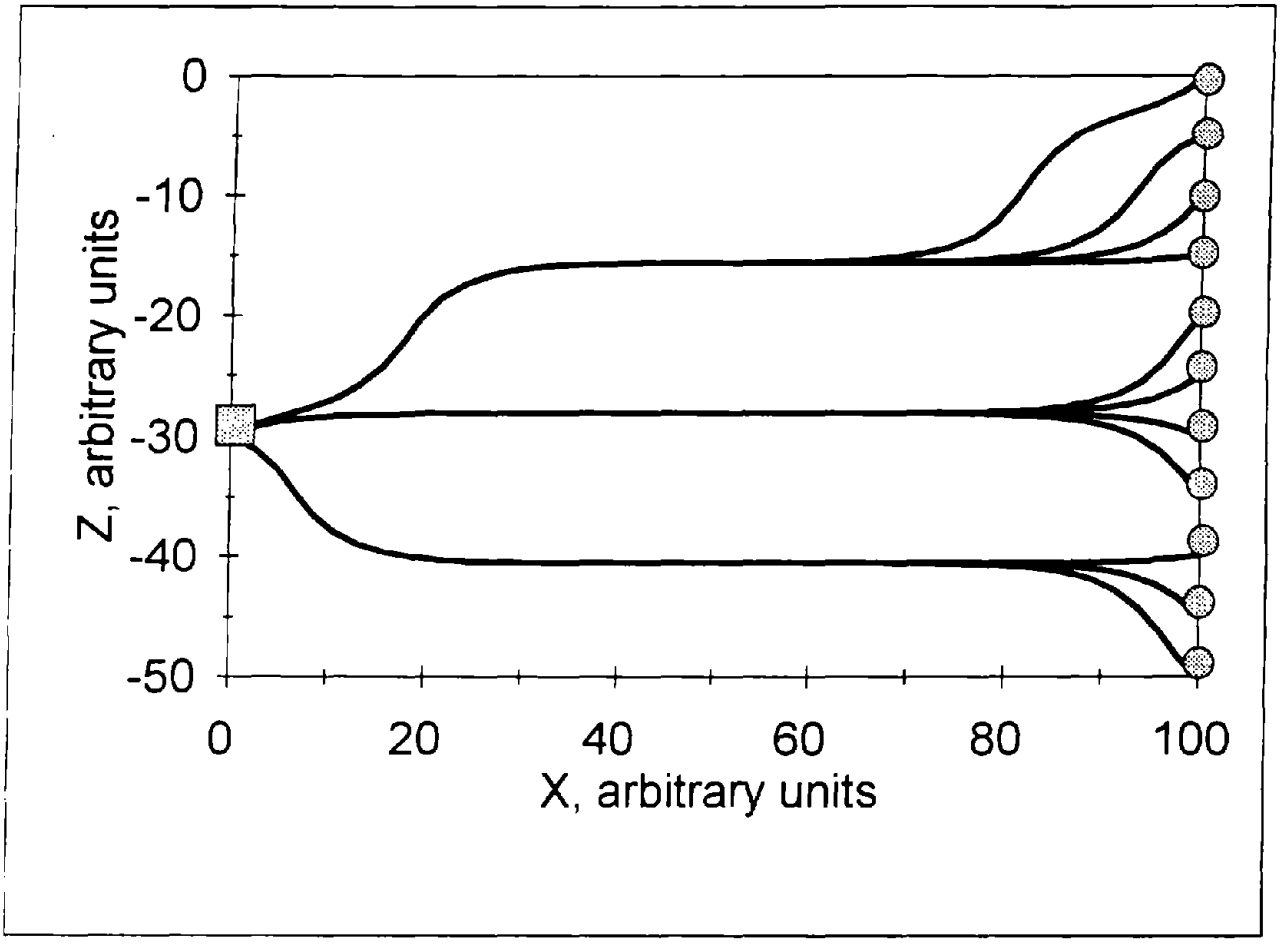

Figure 19. Rays traced using the hybrid approach for the model with velocity varying sinusoidally with depth; four high-velocity layers alternate with slow layers between 0 and 40 depth units.

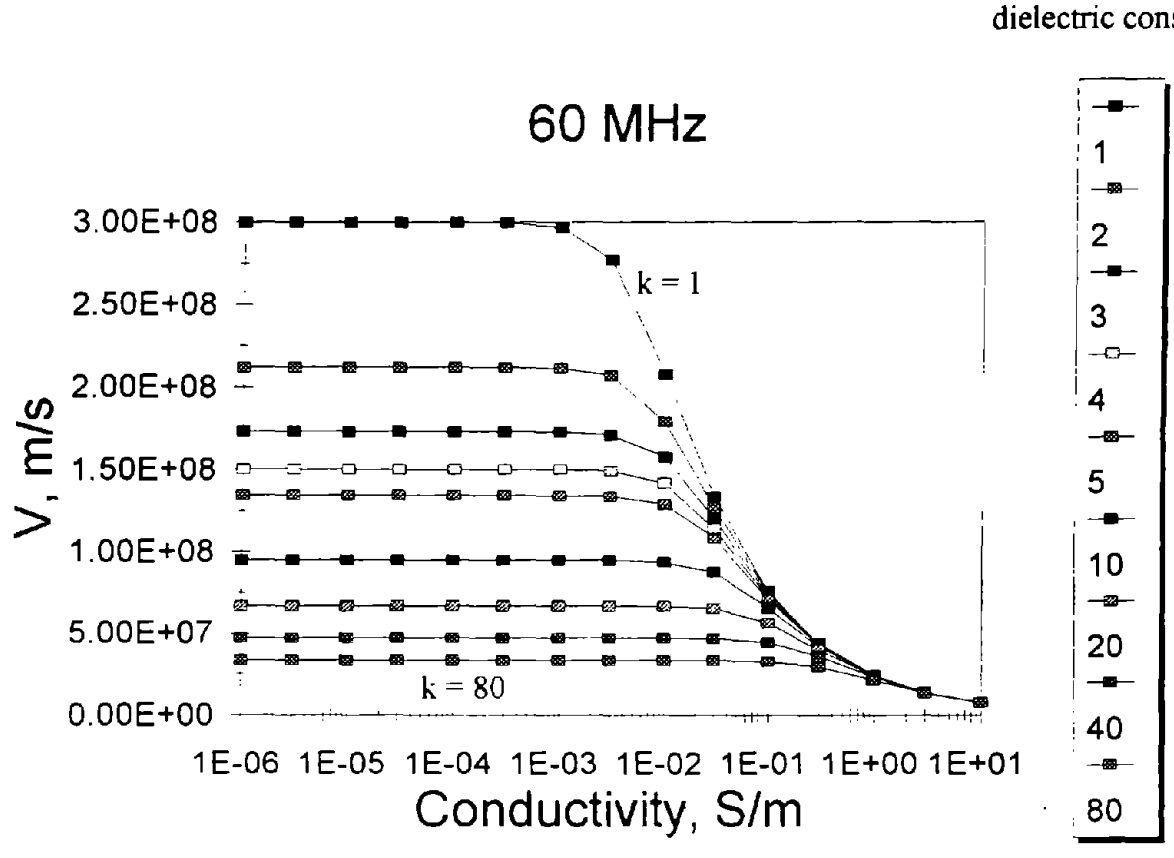

Figure 20. Variation of electromagnetic wave velocity with conductivity and dielectric constant at a frequency of 60 MHz.

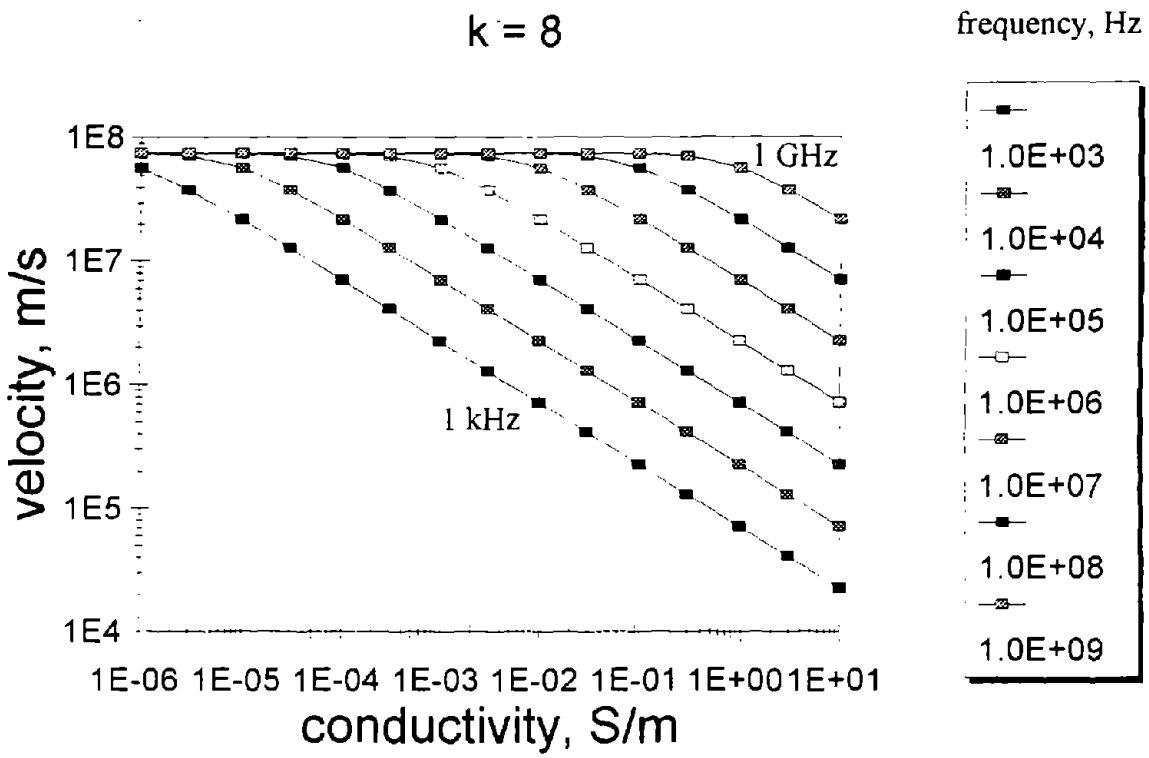

Figure 21. Variation of electromagnetic wave velocity with conductivity and frequency for a dielectric constant of 8.

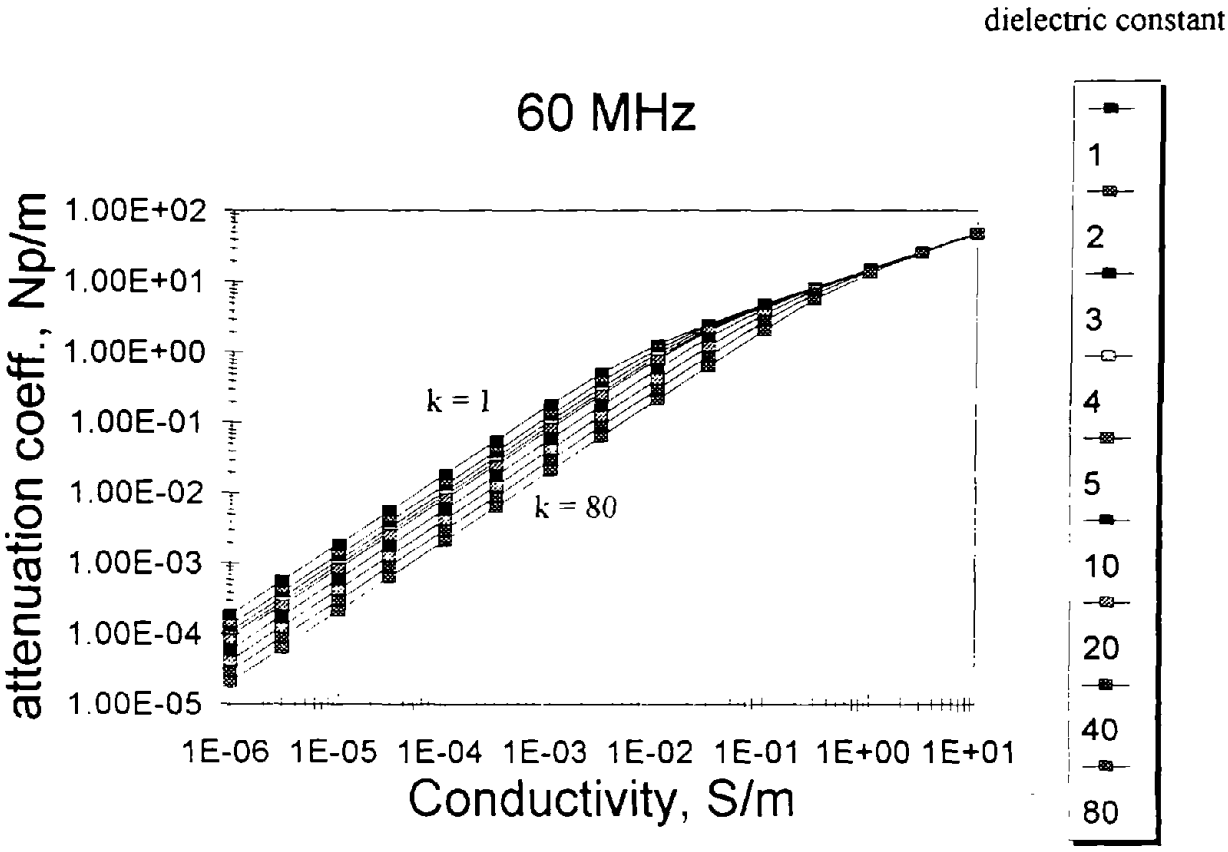

Figure 22. Variation of electromagnetic wave attenuation rate with conductivity and dielectric constant at a frequency of 60 MHz.

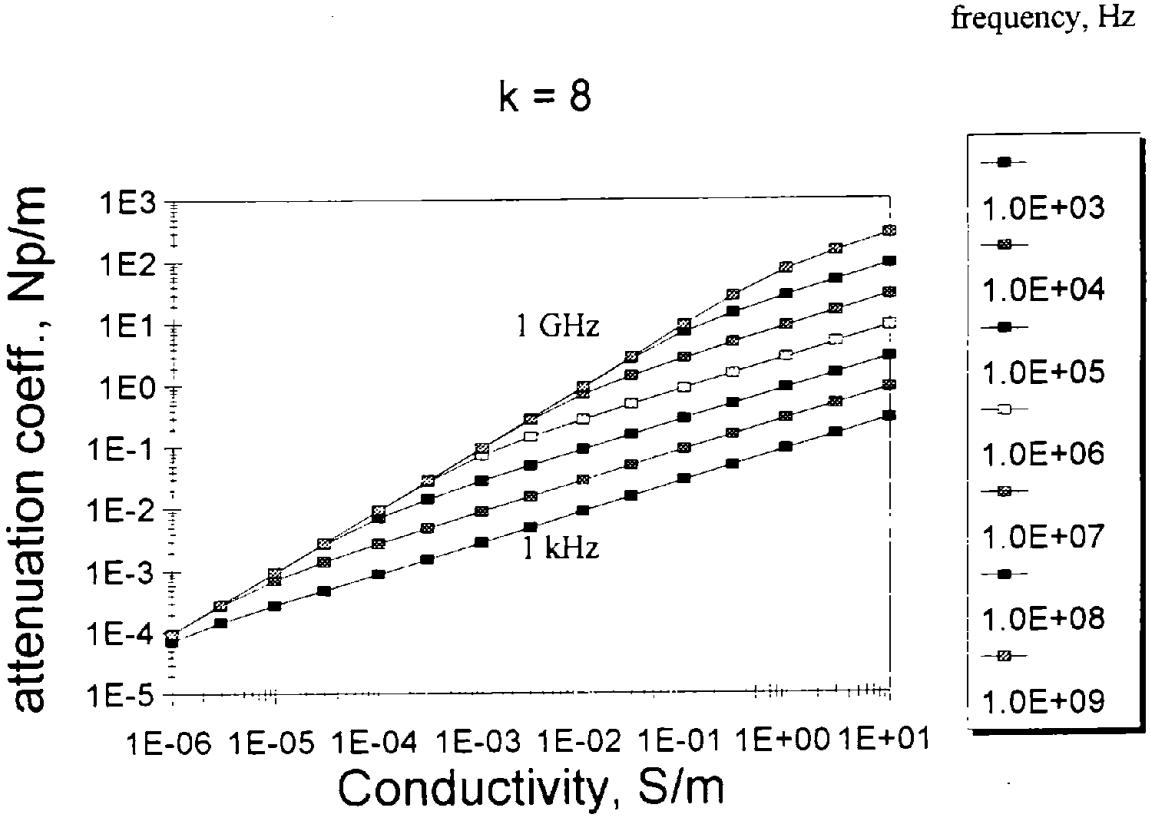

Figure 23. Variation of electromagnetic wave attenuation rate with conductivity and frequency for a dielectric constant of 8.