

Three-dimensional tomographic inversion of combined reflection and refraction seismic traveltimes data

James W. D. Hobro,^{1,*} Satish C. Singh^{2†} and Timothy A. Minshull³

¹Bullard Laboratories, Department of Earth Sciences, University of Cambridge, Madingley Road, Cambridge, CB3 0EZ, UK

²LITHOS, Bullard Laboratories, Department of Earth Sciences, University of Cambridge, Madingley Road, Cambridge, CB3 0EZ, UK

³School of Ocean and Earth Science, University of Southampton, Southampton Oceanography Center, European way, Southampton, 5014 3ZH, UK

Accepted 2002 July 11. Received 2002 July 8; in original form 2000 July 24

SUMMARY

A tomographic inversion method is presented for the determination of 3-D velocity and interface structure from a wide range of body-wave seismic traveltimes data types. It is applicable to refraction, wide-angle reflection, normal-incidence and multichannel seismic data, and is best suited to a combination of these that provides good independent constraints on seismic velocities and interface depths. The inversion process seeks a layer–interface minimum-structure model that is able to explain the given data satisfactorily by inverting to minimize data misfit and model roughness norms simultaneously. This regularized inversion, and the use of smooth functions to describe velocities and depths, allows the highly non-linear tomographic problem to be approximated as a series of linear steps. The inversion process begins by optimizing the fit to the data of a highly-smoothed initial model. In each subsequent step, structure is allowed to develop in the model with successively greater detail evolving until a satisfactory fit to the data is obtained. Parameter uncertainties for the final model are then estimated using an *a posteriori* covariance matrix analysis. Smooth layer–interface models are parametrized using regular grids of velocity and depth nodes from which spline-interpolated interface surfaces and velocity fields are defined. Forward modelling is achieved using ray perturbation theory and a two-point ray tracing method that is optimized for a large number of closely-spaced shot or receiver points. The method may be used to generate 1- and 2-D models (from, for example vertical seismic profile data or 2-D surveys) in which the 3-D geometry of a survey is correctly accounted for. The ability of the method to resolve typical target structures is tested in a synthetic salt dome inversion. From a set of noisy traveltimes data, the model converges quickly to a well-resolved final model from different starting models. The application of this method to real data is demonstrated with a combined 3-D inversion of refraction and reflection data which provide *P*-wave velocity constraints on the methane hydrate stability zone in the Cascadia Margin offshore Vancouver Island.

Key words: inversion, ray tracing, seismic modelling, seismic structure, tomography, traveltimes.

1 INTRODUCTION

The science of estimating the seismic velocity structure of the Earth's crust and upper mantle from seismic traveltimes data has developed rapidly over the last 20 yr. During this time, trial-and-error forward modelling techniques (e.g. Catchings & Mooney 1988; Henry *et al.* 1990), based on repeated modelling by ray-tracing methods (Červený *et al.* 1977) have gradually been superseded by tomographic inversion methods which allow much more objective, and

often simultaneous analysis of large quantities and different types of seismic data. Advances in acquisition technology have enabled large 2-D and 3-D data sets to be produced, and many structures have been imaged using a combination of reflection and refraction seismic techniques. The ray-tracing method is slowly being replaced by more sophisticated full-waveform modelling techniques (e.g. Robertsson *et al.* 1994), but the full potential of ray theory as a tool for tomographic inversion has still not yet been exhausted. Since many seismic data sets are still acquired as 2-D profiles, most tomographic techniques have been developed to produce detailed 2-D images of velocity structure from refraction and/or reflection data (e.g. Spence *et al.* 1985; Huang *et al.* 1986; Firbas 1987; Lutter *et al.* 1990; Zelt

*Now at: WesternGeco, Schlumberger House, Gatwick Airport, UK

†Now at: Laboratoire de Géosciences, IPG Paris, France

& Smith 1992; McCaughey & Singh 1997; Zhang *et al.* 1998). 3-D seismic surveys are now becoming increasingly common, but due to the size and complexity of the problem, the development of 3-D tomographic algorithms (pioneered by Aki & Lee 1976) has so far focused on the inversion of first arrival traveltimes (e.g. Hole 1992; Toomey *et al.* 1994; Zelt & Barton 1998) or direct (refracted) arrivals in the case of earthquake data (e.g. Eberhart-Phillips 1986; Ni *et al.* 1991; Ghose *et al.* 1998). The forward-modelling methods used in these algorithms are designed only to model first arrivals, and are not easily modified to model reflections. As a result, these techniques produce single-layer velocity models, making use of only a small fraction of the data potentially available in this type of survey. It is increasingly common for reflection and refraction data to be available for a single target structure, and these data types have been integrated in several analyses (e.g. Holbrook *et al.* 1994). Simultaneous inversion of reflection and refraction data, however, is still relatively unusual. This approach allows normal-incidence, wide-angle refraction and wide-angle reflection data to be inverted simultaneously, enabling some of the sampling limitations of 2-D wide-angle refraction surveys to be overcome (Wang & Braile 1996; Korenaga *et al.* 2000). The ability of this combined approach to optimize model resolution and reduce velocity–depth trade-offs has been demonstrated by McCaughey & Singh (1997).

In this paper, we present a new method for tomographic inversion of seismic traveltimes in three dimensions. It allows a combination of refraction, wide-angle reflection, and normal-incidence data to be inverted simultaneously to produce a model in which potentially complex structures are represented in the layer–interface formalism. In the conventional approach to tomographic inversion, a single 2-D or 3-D grid of velocity nodes spans the model, from which a velocity field is defined by interpolation (e.g. Ghose *et al.* 1998). We unite this approach with the layer–interface formalism used in traveltimes inversion and velocity analysis by inverting simultaneously for interface depth and layer velocity functions. In addition, smoothing is applied during the inversion by minimizing a

combined function of misfit and model roughness (e.g. Williamson 1990), rather than applying model smoothing separately. This ensures that the solution obtained is correctly optimized to fit the specified misfit and roughness norms, something which cannot be guaranteed if smoothing is applied in between inversion steps. The inversion strategy used is an iterative ‘jumping’ method (Shaw & Orcutt 1985) which selects the smoothest model that provides a satisfactory fit to the data (Fig. 1).

Inversion is achieved as a series of linear steps, in which successively improved models are generated. At each step, the inversion algorithm requires a set of synthetic traveltimes—computed by tracing rays between each source–receiver pair in the current model—and their Fréchet derivatives, the partial derivatives of the traveltimes with respect to each model parameter. These are obtained by using a ray shooting method (e.g. Červený 1987) to map the arrival positions of rays propagating from the source through the model in different directions. Ray shooting is used since it is robust with regard to multipathing, is inherently capable of modelling different ray types (reflected, refracted, converted rays, or any combination of these), and provides an efficient solution for inversion problems of moderate complexity (i.e. up to tens of thousands of model parameters and a similar number of rays). The shooting method may become inefficient when used in highly-complex or densely-parametrized models (e.g. over 100 000 model parameters). In these situations, it may be more appropriate to use an Eikonal-based solver.

Each source–receiver ray path is found by successive linear interpolation, along with its traveltimes and Fréchet derivatives. Ray tracing through the complex smooth media described above is achieved using ray perturbation theory (e.g. Chapman 1985), a semi-analytical method that allows traveltimes and their Fréchet derivatives to be calculated more rapidly than by numerical methods.

This inversion method may be used to generate 1-, 2- or 3-D models from any set of body-wave traveltimes data for which the source–receiver geometry is known. It is therefore applicable to refraction, wide-angle reflection, normal-incidence, multichannel and




Figure 1. A representation of ‘creeping’ and ‘jumping’ strategies (after Shaw & Orcutt 1985) to move from a starting model \mathbf{m}_0 to a final model, which lies within the area of acceptable misfit (shaded region). (a) In the creeping approach, an acceptable misfit model is sought that minimizes the model perturbation $\delta\mathbf{m}_0$. (b) Because of non-linearity, several iterations are normally required. (c) A disadvantage is that the inversion path and final model depend explicitly on the starting model used. (d) A more desirable strategy would produce a final model that is insensitive to small changes in \mathbf{m}_0 . (e) The jumping strategy seeks a final model that minimizes a norm relative to an absolute origin \mathbf{O} , removing the explicit dependence on the starting model. (f) An iterative form of this approach may be used to solve non-linear problems by moving from smooth to rough best-fit models until an acceptable misfit is obtained.

vertical seismic profile data. It is best suited to a combination of these that provides good independent constraints on seismic velocities and interface depths.

In the following sections, the method is described in detail. Its properties are then examined in a synthetic example, and finally, it is applied to a real 3-D traveltimes data set.

2 INVERSION APPROACH

Traveltimes inversion is a flexible technique that has a wide range of applications, from the simple determination of seismic velocity gradients to the more sophisticated tomographic modelling of complex geological structures. Analysis of crustal-scale wide-angle seismic data by picking traveltimes has generally led to the visualization of the Earth as a series of layers with smoothly-varying (or constant) velocity gradients, separated by interfaces, or discontinuities in seismic velocity. These structures are then often extended to form two- or 3-D models which are able to explain the data available, often through a combination of trial-and-error forward modelling and inversion of a few selected parameters. A data set may potentially be explained by many different possible models. Most such models are recognizably non-physical, and a means of selecting preferred models must be found. When models are generated by extensive user input, a level of subjectivity is introduced into this process, since the complex relationship between features in the model and features in the data may be subtle and difficult to visualize.

Although target structures typically contain seismic velocity variations at many different spatial frequencies, the band-limited nature of seismic data and the additional resolution loss that results from the picking of traveltimes justify the selection of smooth models in order to overcome this problem of uniqueness. We propose that, in accordance with Occam's razor (Constable *et al.* 1987), and in recognition of the limitations in acquisition and forward-modelling methods, it is useful to seek the smoothest tomographic models that are able to explain the data within the layer–interface formalism outlined above. In order to achieve this, the variable smoothing method (e.g. Williamson 1990) is adopted: a very smooth model is first optimized to produce the best possible fit to the data, then structure is allowed to emerge in the model with successively greater detail in order to produce an improved fit until the data have been explained satisfactorily. This approach, and the use of smoothly-varying interface and velocity structures allows the highly non-linear tomographic problem of inverting for seismic reflection and refraction traveltimes to be approximated effectively as a series of linear steps. It also produces a model that contains the minimum degree of structure that is required to fit the data.

2.1 Model parametrization

An important feature of any traveltimes inversion method is its approach to model parametrization. Simple models based on a few parameters may be well-constrained by the data available, but are generally unable to describe the complexity of most target structures. Some inversion schemes (e.g. Zelt & Smith 1992) extend this approach to deal with complex structures by allowing the user to change the model parametrization density in order to match structure observed in the data. In seeking to automate this process, we use a density of parametrization that is sufficient to explain the most detailed structural features expected. This approach of over-parametrization requires smoothing regularization to stabilize the inversion process, and to allow the degree of model structure to be controlled.

Models are defined as a series of layers separated by interfaces. Within each layer a potentially complicated seismic velocity field is modelled as a function of position. The interfaces represent discontinuities in seismic velocity, at which reflections and refractions may occur, and whose depth is modelled as a function of lateral position. In both cases, the functions are continuous and smooth (i.e. their first spatial derivatives are continuous) and are derived from a regular grid of velocity or depth parameters using quadratic or cubic B-splines (Buchanan & Turner 1992, Chapter 5). Structure with variation up to a level equal to the grid spacing in each layer or interface may be modelled without introducing aliasing artefacts, ensuring that the total number of parameters is kept to a minimum. A fixed ratio of *P*- to *S*-wave velocity is defined for each layer.

The extent of each layer within the model is defined by the positions of the two interfaces that bound it. The grid that maps the layer is fixed however, and covers a cuboid which must encompass the maximum expected extent of the layer as the model evolves during the inversion process. Layers and interfaces may if required be defined to span only part of the model so that unconstrained areas of the model are left undefined, reducing the number of redundant inversion parameters.

Interfaces may be joined together by mapping equivalent parameters on adjacent interfaces to a single inversion parameter, allowing layers to 'pinch out' in some regions of the model. Similarly, adjacent layers may be bound together by specifying regions in which smoothing constraints operate across an interface.

2.2 The forward problem

The ray approximation provides an efficient means of obtaining synthetic traveltimes from a velocity model, and is able to resolve considerable detail in tomographic imaging. Direct analytical computation of ray paths and traveltimes is not possible through the complex smooth velocity fields described above, but the use of ray perturbation theory allows a semi-analytical solution to be obtained more rapidly than by numerical methods. We adopt a method based on the approach of Virieux & Farra (1991) which is well-adapted to our parametrization and simplifies the calculation of traveltimes and their Fréchet derivatives (McCaughey & Singh 1997). The use of smooth models ensures that ray paths through the model, and therefore the calculated synthetic traveltimes, change in a predictable manner in response to perturbations to the model parameter values. This greatly enhances the stability of the inversion process.

The geometry of a data set is described by defining a set of 'sources' and 'receivers', all of which are points within the model. Rays are traced from each source, through the model to a list of receivers, a process which is optimized for a small number of sources and a large number of receivers. Sources may be located anywhere in the model. Receivers must either be located on a horizontal plane (marine geometry) or on the first interface in the model, which represents the Earth's surface in land geometry. The principle of traveltimes reciprocity allows this formalism to be adapted to a wide range of different experimental geometries. In a wide-angle marine experiment, sources may be used to represent ocean-bottom instruments, and receivers to represent shot points. For a controlled-source experiment on land, where the number of seismic sources and instruments may be comparable, either reciprocal arrangement may be used. Travel-time data are divided into 'phases', which describe the sequence of layers through which each ray passes between source and receiver points.

The forward problem—the determination of traveltimes and Fréchet derivatives for each source–receiver pair—is solved in two

stages. A 3-D ‘fan’ of rays is first propagated from each source, for each ray phase, through the model to the receiver surface (the shooting method). These rays are used to map the model by exploring all possible routes from the source to the surface containing the receivers. At this stage, ray tracing is extremely efficient as the only information required for each ray is its point of arrival at the receiver surface, so the computation of traveltimes and Fréchet derivatives is unnecessary. The response of the model to these rays may be represented by the function

$$\mathbf{x} = \mathbf{f}(\Theta) \quad (1)$$

where $\mathbf{x} = (x, y)$ is the point of arrival at the receiver surface, and $\Theta = (\theta, \phi)$ specifies the direction in which the ray is travelling as it leaves the source. This is a 2-D function of a 2-D variable which maps all model structure covered by the ‘fan’ of rays, including interface and velocity features, and the topography of the receiver surface. The function is divided into segments in which \mathbf{x} is a monotonic function of Θ (Fig. 2), and the maximum and minimum values of x & y in each segment are recorded. For a complex 3-D model there may be dozens or even hundreds of segments. The extent and density of the ‘fan’ must be sufficient to map all structure in the model that may be sampled by the data.

In the second stage of the forward problem—the two-point problem—each source–receiver pair is considered in turn. By comparing the receiver position $(x_{\text{rec}}, y_{\text{rec}})$ with the maximum and minimum values of x & y for each segment, the segments in which a solution may exist are quickly selected, and within these a solution is sought to within a specified tolerance by successive linear interpolation of previously-traced ray data. At this stage, traveltimes and Fréchet derivatives are calculated for each ray. In many cases,




Figure 2. A map of the domain of the emergence function $\mathbf{x} = \mathbf{f}(\Theta)$ for reflections from an interface in a relatively simple 3-D model. This example is sampled at 1000×1000 equally-spaced points in the domain by shooting rays at the take-off angle $\Theta = (\theta, \phi)$ from a single point. Most rays reach the model surface, yielding an emergence position $\mathbf{x} = (x, y)$. The turning points in the function $\mathbf{x} = \mathbf{f}(\Theta)$ (marked in black) divide the domain into segments in which the function is monotonic. The grey areas are regions in which rays do not reach the model surface and the function is therefore undefined.

consecutive receiver positions are closely-spaced, so a ray solution may be found on the first attempt by using data from previously-traced rays. In marine experiments, where shot positions are densely-spaced in lines, tests have shown that the ray-tracing overhead for this method is around 10 per cent (i.e. the number of rays that are traced is 10 per cent greater than the number of receivers). Where multiple arrivals exist within the same phase for an individual source–receiver pair, the solution with the smallest traveltime is selected. This method has the advantage over many other two-point techniques (e.g. Moser 1991) that where multiple solutions exist, they are identified and a consistent selection is made.

2.3 The inverse problem

In order to describe the mathematics of the inverse problem, we define a general model vector \mathbf{m} , an ordered set of all model parameters included in the inversion, and a general traveltime vector \mathbf{t} , an ordered set of real or synthetic traveltimes. At each step the inversion process operates on the current model vector \mathbf{m}_i (for the i th step) producing a model update vector $\delta\mathbf{m}_i$ from the real traveltime data (\mathbf{t}_{real}), the set of synthetic traveltimes derived from the current model (\mathbf{t}_i), and the matrix of Fréchet derivatives associated with the synthetic traveltimes (\mathbf{A}_i). The traveltime residual vector,

$$\mathbf{r}_i = \mathbf{t}_{\text{real}} - \mathbf{t}_i \quad (2)$$

is the subject of the optimization, which is based upon an approximation in which a perturbation in the traveltime data vector $\delta\mathbf{t}_i$ is related to a perturbation in the current model ($\delta\mathbf{m}_i$) by the linear equation

$$\delta\mathbf{t}_i = \mathbf{A}_i \delta\mathbf{m}_i, \quad (3)$$

where

$$[\mathbf{A}_i]_{jk} = \frac{\partial t_j}{\partial m_k} \quad (4)$$

are the elements in the matrix of Fréchet derivatives. This is simply the first term of a Taylor expansion and it is valid only for small changes in the model, within what may be defined as a ‘region of linearity’ surrounding the current model in the space of all possible models. This relation is justified by McCaughey & Singh (1997), who derived expressions for the terms of \mathbf{A}_i using ray perturbation theory. At each inversion step, an improved fit to the data is sought by perturbing the vector \mathbf{m}_i by an amount $\delta\mathbf{m}_i$. The new model,

$$\mathbf{m}_{i+1} = \mathbf{m}_i + \delta\mathbf{m}_i \quad (5)$$

will in general only provide a better fit if it lies within the region of linearity surrounding \mathbf{m}_i .

Eq. (3) cannot be solved directly, as this solution would generally lie outside the region of linearity. In addition, we wish to apply smoothing constraints to the solution. One of the most useful and flexible ways of formulating the problem is as a least-squares optimization, which is well-suited to this inversion problem for the following reasons:

- (i) the largest data residuals, which are generally resolved by large-scale changes in the model, are accounted for preferentially during inversion, and
- (ii) this formulation allows parameter uncertainties to be obtained for the final model relatively easily by an *a posteriori* covariance matrix analysis.

A general least-squares objective function may take the form

$$\begin{aligned} \mathcal{F}(\delta \mathbf{m}_i) = & \|\mathbf{r}_i - \mathbf{A}_i \delta \mathbf{m}_i\|_D^2 \\ & + \lambda_\alpha \|\mathbf{m}_i + \delta \mathbf{m}_i - \mathbf{m}_\alpha\|_{M_\alpha}^2 \\ & + \lambda_\beta \|\mathbf{m}_i + \delta \mathbf{m}_i - \mathbf{m}_\beta\|_{M_\beta}^2 \\ & + \dots \end{aligned} \quad (6)$$

Here, $\mathcal{F}(\delta \mathbf{m}_i)$ is a misfit function and the subject of the optimization. It contains a single term which measures data misfit and a series of terms which measure properties of the model. $\|\cdot\|_D$ and $\|\cdot\|_{M_\alpha}$, $\|\cdot\|_{M_\beta}$, ... are weighted l_2 norms, and λ_α , λ_β , ... are scalars which control the emphasis given to each of these terms during the inversion process. This formulation may be used to constrain the inversion in a variety of ways, for example by preventing the model from changing too much ($\mathbf{m}_\alpha = \mathbf{m}_i$; the damped least-squares method), by introducing *a priori* information in a particular region of the model, or by constraining the model to be flat or smooth. We adopt the form

$$\mathcal{F}(\delta \mathbf{m}_i) = \|\mathbf{r}_i - \mathbf{A}_i \delta \mathbf{m}_i\|_D^2 + \lambda \|\mathbf{m}_i + \delta \mathbf{m}_i\|_M^2 \quad (7)$$

after Tikhonov & Arsenin (1977) where $\|\cdot\|_D$ weights each residual according to the uncertainty in the equivalent traveltimes datum, and $\|\cdot\|_M$ measures the roughness of each layer and interface in the evolving model. Note that the model is not constrained to be close to any particular model (e.g. the starting model, or the result of the previous step), ensuring that the model evolution is not strongly determined by the starting model. The norms $\|\cdot\|_D$ and $\|\cdot\|_M$ contain a variety of weightings and normalization factors. As some of the model parameters are depths and others velocities, their units and therefore their magnitudes may be quite different. This would result in a bias towards changing one type of parameter during the inversion if no normalization were applied. To remedy this, the model parameters of each type are normalized by the factor

$$n = \sqrt{\sum_{\text{all } j,k} [\mathbf{A}_i]_{jk}^2} \quad (8)$$

where the sum is over all Fréchet derivatives of that parameter type, and the equivalent Fréchet derivatives are normalized by the reciprocal of n . This balances the gradient of the function $\mathcal{F}(\delta \mathbf{m}_i)$ evenly between the two parameter types whilst ensuring that the product $\mathbf{A}_i \delta \mathbf{m}_i$ is unchanged. The two norms each contain a single normalization constant which is designed to bring both terms close to unity. This prevents the number of data and the model geometry from having a strong effect on the values of λ required to achieve a stable inversion. Extra weightings may also be applied within $\|\cdot\|_M$ to adjust the relative smoothing constraints between interfaces and the lateral and vertical velocity fields. A weighting of $1/n_{\text{total}}$ is applied to $\|\cdot\|_D$ where n_{total} is the total number of traveltimes data available. The vector \mathbf{r} contains n_{res} elements where $n_{\text{res}} \leq n_{\text{total}}$, as it may not have been possible to obtain residuals for all the given data. The weighting applied to $\|\cdot\|_D$ is calculated from n_{total} and not n_{res} in order to stabilize the inversion process—if an inversion step causes n_{res} to decrease, the size of $\|\cdot\|_D$ also decreases and the inversion will tend to smooth the model, preventing it from changing to provide a good fit to only a small subset of the data.

The roughness term $\|\mathbf{m}\|_M^2$ for a model \mathbf{m} contains an approximation to a series of integrals which measure the roughness of each interface and velocity layer in the model. In the 3-D case, the two components are

$$\|\mathbf{m}\|_{M_1}^2 = S \int_S \left(\frac{\partial^2 z}{\partial x^2} \right)^2 + \left(\frac{\partial^2 z}{\partial y^2} \right)^2 + 2 \left(\frac{\partial^2 z}{\partial x \partial y} \right)^2 dS \quad (9)$$

for interfaces where S is the area of the interface and $z(x, y)$ is the function describing the interface, and

$$\begin{aligned} \|\mathbf{m}\|_{M_V}^2 = & V^{1/3} \int_V \left(\frac{\partial^2 v}{\partial x^2} \right)^2 + \left(\frac{\partial^2 v}{\partial y^2} \right)^2 + \left(\frac{\partial^2 v}{\partial z^2} \right)^2 \\ & + 2 \left(\frac{\partial^2 v}{\partial x \partial y} \right)^2 + 2 \left(\frac{\partial^2 v}{\partial y \partial z} \right)^2 + 2 \left(\frac{\partial^2 v}{\partial z \partial x} \right)^2 dV \end{aligned} \quad (10)$$

for velocity layers where V is the volume of the velocity layer and $v(x, y, z)$ is the velocity function. The integrands are approximated using the finite-difference operators

$$\left(\frac{\partial^2 m}{\partial x^2} \right)_{ij} = \frac{1}{(\Delta x)^2} (m_{i-1,j} - 2m_{i,j} + m_{i+1,j}) \quad \text{and} \quad (11)$$

$$\begin{aligned} \left(\frac{\partial^2 m}{\partial x \partial y} \right)_{ij} = & \frac{1}{\Delta x \Delta y} (m_{i-1,j-1} - m_{i+1,j-1} \\ & - m_{i-1,j+1} + m_{i+1,j+1}) \end{aligned} \quad (12)$$

where m_{ij} is the model parameter at grid index (i, j) , and Δx & Δy represent the parameter separation along the x - and y -axes. A set of equivalent first derivative terms which span the lateral model edges are also included to penalize lateral gradients in interface depth and velocity, whilst allowing vertical velocity gradients to develop unrestrained.

A complete minimization of this misfit function is not desirable, as it would in general take the model update vector outside the region of linearity surrounding the current model. Instead, we require a local method that will move towards the minimum value of $\mathcal{F}(\delta \mathbf{m}_i)$ whilst remaining within the region of linearity. We use the conjugate gradient method (Hestenes & Stiefel 1952), which has been used successfully to solve large tomographic problems in seismology (Nolet 1987), and is easily adapted to take advantage of the sparseness of the matrix \mathbf{A}_i . The number of iterations required by this method to reach the minimum of $\mathcal{F}(\delta \mathbf{m}_i)$ is equal to the number of dimensions in the parameter space—in this case, the number of model parameters for which we are inverting. The minimization process may be halted after a smaller number of iterations (Fig. 3), yielding a solution that is closer to the minimum, but which still lies within the region of linearity.

A series of these inversion steps are taken at a constant value of λ , halting the minimization process at successively later stages, and finishing with a complete minimization. During this process, the model converges to a solution that is optimized for the given value of λ . Although each individual step must be small, the sum of all these steps can change the model dramatically if such a change is required in order to fit the data.




Figure 3. Convergence of the two objective function terms (data misfit and model roughness) during a single inversion iteration: in this example, the model changes significantly in order to produce an improved fit to the data. Most of this convergence occurs during the first 10 per cent of the full optimization process.

In order to produce a minimum-structure model, inversion is first performed as described above at a value of λ chosen to be high enough not to allow any detailed structure to appear in the model. During this first step, bulk velocities, velocity gradients and the average depths of interfaces are obtained, and the model is optimized to fit the data at this level of ‘zero’ structure. The value of λ is then decreased in steps, and an optimized model solution is produced at each step. This process is repeated until a satisfactory fit is obtained. A crude estimate of the quality of fit may be obtained by evaluating the normalized χ^2 parameter,

$$\chi^2 = \frac{1}{n_{\text{res}}} \sum_{j=1}^{n_{\text{res}}} \left(\frac{r_j}{\sigma_j} \right)^2 \quad (13)$$

where r_j is the element of \mathbf{r} corresponding to the j th traveltime datum, and σ_j is the uncertainty in that datum. If the uncertainties are well-estimated, uncorrelated and follow a Gaussian distribution then a satisfactory fit is obtained on average across the model when $\chi^2 = 1$. In reality, variations in the quality of fit across the model and deviation of uncertainties from a well-estimated normal distribution often mean that χ^2 is suitable only as an approximate gauge of the quality of fit, which should be estimated by examining an overview of the traveltime residuals for the entire data set.

In principle, a satisfactory fit is achieved when all structural features in the real data have been modelled to within the estimated picking error by synthetics, but before the fit to the real data is detailed enough to model the visible noise. In practice, a more precise fit to the real data will be obtained with each inversion step, and a judgement of whether this condition has been met must be made at each step, based on a visual inspection of residuals for the whole data set. If the spatial frequencies of the data signal and noise components are well separated, the decision may be straightforward. If this separation is less well defined, then sharp features in the data signal (e.g. faults in interfaces) may be partially smoothed out in the final model. In more complex cases, the point at which to stop the inversion is not so easily defined, and several possible final models containing different levels of detail should be considered.

2.4 Uncertainty analysis

The estimation of uncertainties in model parameters is crucial to the interpretation of any velocity model. Under the linearizing approximations of Section 2.3, uncertainties in the values of the model parameters in the final model may be estimated by examining the *a posteriori* covariance matrix for the inversion (Tarantola 1987),

$$\mathbf{C}'_{\mathbf{M}} = \mathbf{H}^{-1} \quad (14)$$

where \mathbf{H} is the Hessian matrix, which measures the curvature of the misfit function. If the norms $\|\cdot\|_{\mathbf{D}}$ and $\|\cdot\|_{\mathbf{M}}$ in eq. (7) are expanded in terms of their corresponding covariance matrices $\mathbf{C}_{\mathbf{D}}$ and $\mathbf{C}_{\mathbf{M}}$, the Hessian matrix may be written

$$\mathbf{H} = 2(\mathbf{A}^T \mathbf{C}_{\mathbf{D}}^{-1} \mathbf{A} + \lambda \mathbf{C}_{\mathbf{M}}^{-1}). \quad (15)$$

Computing the *a posteriori* covariance matrix therefore involves the calculation and inversion of the Hessian matrix, neither of which were necessary to minimize the misfit function by the conjugate gradient method. In the examples presented in this paper, the Hessian was calculated and inverted in its full, uncompressed form since its irregular sparse nature did not allow an advantage to be gained by treating it as sparse during inversion. The inversion routine DSYTRI from the public domain library LAPACK was used. This routine is optimized for symmetric matrices, and inverted a $12\,000 \times 12\,000$

Hessian in around 2 hours on a single node of a Silicon Graphics Origin 2000 (MIPS R12000 300 MHz CPU).

The diagonal elements of $\mathbf{C}'_{\mathbf{M}}$ measure the variance in each model parameter, and the other elements contain information on correlations between parameters. Measures of the standard error σ_i for the i th parameter, and the correlation ρ_{ij} between parameters are given by (Tarantola 1987)

$$\sigma_i = \sqrt{C'_{ii}} \quad \text{and} \quad \rho_{ij} = \frac{C'_{ij}}{\sqrt{C'_{ii} C'_{jj}}}. \quad (16)$$

These uncertainties and correlations are valid for small perturbations around the final model in the model space of a set of discrete velocity and depth parameters. They do not take into account the quality of fit of any models which differ significantly from the final model. A large positive ($\rho_{ij} \approx 1$) or negative ($\rho_{ij} \approx -1$) correlation indicates that the two parameters are not independently constrained.

2.5 Data preparation

The manner in which data are prepared for traveltime inversion has important implications for the construction of models from these data. Any errors of measurement or misinterpretation in acquisition, processing or traveltime picking will propagate into models as spurious model structure or as systematic errors in model parameter values. In order to eliminate, or at least minimize the impact of these errors, a number of precautions should be taken when processing and picking the data. Systematic errors in picked traveltimes are not accounted for by the least-squares inversion formulation used in this method, so any static shifts in seismic sections should be corrected before traveltimes are picked. The seismic data that will be modelled must then be divided into ‘phases’—horizons for which an energy path (i.e. the sequence of reflections and refractions through interfaces and layers from source to receiver positions) has been identified with confidence. In some cases, a trial-and-error approach may be necessary to distinguish between different phase assignments by examining models produced by different interpretations of the data in order to select a preferred interpretation.

Once all the phases that are to be used have been identified, traveltimes in each phase may be picked. Each traveltime is assigned an individual uncertainty value, allowing the uncertainties due to changing signal-to-noise ratio or the variation of waveforms with offset to be accounted for. The least-squares optimization method uses a Gaussian distribution function to model noise in the traveltimes, and non-Gaussian components of this noise will be interpreted as structure within the data set. It is therefore important to ensure that the traveltime picks do not contain ‘outliers’ or spurious values lying a long way outside the specified uncertainty range that may be erroneously modelled during inversion. The uncertainty values given should define a region within which approximately 66 per cent of the traveltime picks fall.

3 TESTS ON SYNTHETIC DATA

It is important to test the inversion algorithm on a synthetic data set for several reasons:

- (i) to assess its ability to recover a known realistic structure from a set of noisy traveltime data,
- (ii) to evaluate any artefacts that may be introduced during inversion,

(iii) to assess the effect of different starting models upon the inversion,

(iv) to study the convergence of the inversion path and

(v) to evaluate the performance of the uncertainty analysis.

A practical inversion algorithm must be able to resolve structures which typically form targets in seismic surveys. In order to test the effectiveness of this algorithm in resolving a realistic structure, a synthetic test was devised using a model in which a salt dome intrusion penetrated a sedimentary sequence containing a bright reflector. The salt dome model used here contained a velocity gradient representing a sedimentary sequence approximately 2 km thick, bounded

at its base by a gently sloping reflector, interrupted by a salt dome intrusion of high seismic velocity. Two sections through the model are plotted in Figs 4(a) and (b). A synthetic data set was generated from this model using the forward-modelling method described in Section 2.2. A grid of instruments ('sources') was located at the model surface, and a set of intersecting lines was constructed along which shot points ('receivers') were placed at regular intervals. This geometry is illustrated in Fig. 4(c). No reflections were allowed from the interface within the area covered by the salt dome (in the region $-6 < x < 1$, $-6 < y < 1$), simulating the disruption of the reflector by the intrusion of the dome, and preventing the reflection data in this experiment from providing unrealistically good constraints




Figure 4. Model, survey geometry and data for the synthetic salt dome experiment: (a) horizontal and (b) vertical cross-sections through the model used to generate synthetic data—the dotted section of the interface indicates the region in which no reflections were permitted to simulate disruption of the reflector by the dome; (c) geometry of the synthetic experiment, with 16 surface instruments (circles) and 290 shot points (dots). The data for a single instrument (solid circle) from all shots are plotted for (d) refracted and (e) reflected arrivals. The three refracted phases which pass through the salt dome are marked and correlated with their corresponding shot lines in (c).

within the dome. Sets of normal-incidence, wide-angle reflection and wide-angle refraction traveltimes data were generated from this model and experimental geometry between all source–receiver pairs for which solutions existed. Pseudo-random Gaussian noise with a constant standard deviation of 10 ms was added to all the traveltimes data. The wide-angle data for a single instrument (containing all shots) are plotted in Figs 4(d) and (e), with the traveltimes reduced at constant velocity. In this plot, the sign for the offset was obtained from the x -coordinate of the shot point relative to the source, so the data characteristics at positive offset are influenced only by the constant gradient background velocity, while those at negative offset also contain features caused by the presence of the salt dome. The refracted arrival curve contains three distinct phases produced by sampling the salt dome velocity structure along the three shot lines marked in Fig. 4(c). The effect of the dome on the reflected arrivals is less marked, and is manifest as a spread towards earlier traveltimes at negative offsets.

These data were inverted, employing the same velocity grid used to define the model ($10 \times 10 \times 10$ nodes) and a 10×10 node interface grid. A constant gradient, horizontal interface starting model was chosen in which the gradient and interface depth were both different from those in the real model. No *a priori* information was included in the form of, for example, weighting factors between the smoothing applied to velocity and interface functions, or between the lateral and vertical velocity smoothing. The inversion proceeded in 8 steps. Within each step, 4 iterations of forward-modelling and inversion were performed with the regularization strength held at a constant value. The regularization strength was then reduced for the following inversion step. A statistical analysis of the model evolution during inversion is shown in Fig. 6, and a series of vertical sections through the model as it evolved during inversion is shown in Fig. 5. The sections shown all follow the same slice as the vertical section in Fig. 4(b).

This model sequence bears close examination. During the first stage of the inversion (models a and b; the 1-D optimization), the average velocity, velocity gradient and interface depth all converge to values which minimize the data misfit. At this stage no lateral variation is allowed and the velocity function with depth is constrained by regularization. An optimal 1-D model is obtained and the misfit considerably reduced. During the early stages of the 3-D optimization (models c and d), the features of the model emerge but are highly smoothed due to the strength of the regularization. At this stage the interface is well-positioned and its slope is correct, and the salt dome is beginning to emerge in the correct position. As the inversion proceeds (models e–h), the salt dome becomes increasingly well-defined. In the final model, for which $\chi^2 = 2$, a satisfactory fit to the data, based on an examination of the traveltimes residuals throughout data set, was judged to be close to being achieved. In fact, a stable solution with $\chi^2 = 1$ could not be found in this case—this was attributed to very high velocity gradients present within the salt dome, which could not be modelled satisfactorily using a single-layer model.

Residuals for a single section of wide-angle refraction and reflection data and for normal-incidence data are plotted in Fig. 5, for an early model and the final model (steps 4 and 8). The early residuals are small at positive offsets, where the velocity and interface functions have correctly been determined, and larger at negative offsets, in the region of the salt dome, especially in the refracted arrivals. In the final model the residuals in this region are much improved, and the importance of the refraction data in determining the velocity structure of the dome may be seen by comparing plots (i) and (l) in Fig. 5.

Sections through the real and final models for this inversion are plotted together in Fig. 7 alongside a plot of uncertainties in the final velocity function, estimated using the method described in Section 2.4. A comparison between the discrepancy in the final model (the velocity in the real model minus the velocity in the final model) and the estimated uncertainty along a line through the salt dome is plotted in Fig. 8. The plot of uncertainties indicates, as expected, that the edges and deeper corners of the model are poorly constrained. There is also a region within the salt dome, in which the high velocity gradients cause seismic energy to be deflected away and in which no reflections were produced, with the result that the interior of the dome is poorly sampled. Fig. 8 indicates that the uncertainty in this area correctly indicates the extent to which velocities have been underestimated relative to the discrepancy in the surrounding area. At the model edges, where uncertainties are large, the actual discrepancy is small because the real model was smooth in these areas and an identical model has been constructed as a result of the smoothing regularization. Interestingly, the uncertainty analysis consistently overestimates the actual discrepancy in the final model, although it is a good indication of the relative uncertainties across the model. It is not clear whether such overestimation will always occur.

The salt dome inversion illustrates the most significant artefact introduced into models produced by the regularized least-squares inversion method. The later models in the sequence plotted in Fig. 5 all contain small fluctuations about an otherwise smooth model in the velocity contours away from the salt dome, and in the interface. Since the data constrain these areas to be completely smooth and the regularization applies additional smoothness constraints, the presence of these fluctuations requires some explanation. Their presence is due to an effect that resembles *Gibbs' phenomenon*, or the behaviour of a damped simple harmonic system in response to a sudden change. Gibbs' phenomenon is the 'ringing' effect that may be produced when high-frequency components are removed from some waveforms (e.g. a time-series square wave, which is constructed from an infinite series of sinusoidal components). The similarity with the phenomenon described here occurs because the smoothing process suppresses the highest frequency components of the velocity function. This 'ringing' is produced by the regularized least-squares algorithm as it attempts to model a structure in which the velocity gradients or gradients in interface depth are very high. The model produced is still the smoothest model that is able to fit the data to the selected level in a least-squares sense.

In the salt dome model (Fig. 5), the ringing emerges as the structure of the salt dome begins to form (step 5) and increases as the dome boundaries are defined (step 6), then diminishes as the velocity structure within the dome is modelled. The ringing in the interface is introduced as a result of the velocity–depth trade-off which links the evolution of velocity and depth features. This ringing effect is illustrated for a simple 1-D velocity inversion in Fig. 9.

The effect of different starting models on the inversion process was tested by running the salt dome inversion again with a radically different starting model. During the initial 1-D optimization, the model obtained was indistinguishable from the optimized 1-D model from the previous inversion (compare Figs 10b and e). The 3-D optimization then proceeded in a similar manner for both inversions, converging on almost identical models. A comparison between model cross-sections at each stage of these two inversions is given in Fig. 10. In this synthetic inversion, in which a relatively simple target structure was well-constrained by the data, the result obtained was independent of the starting model used.

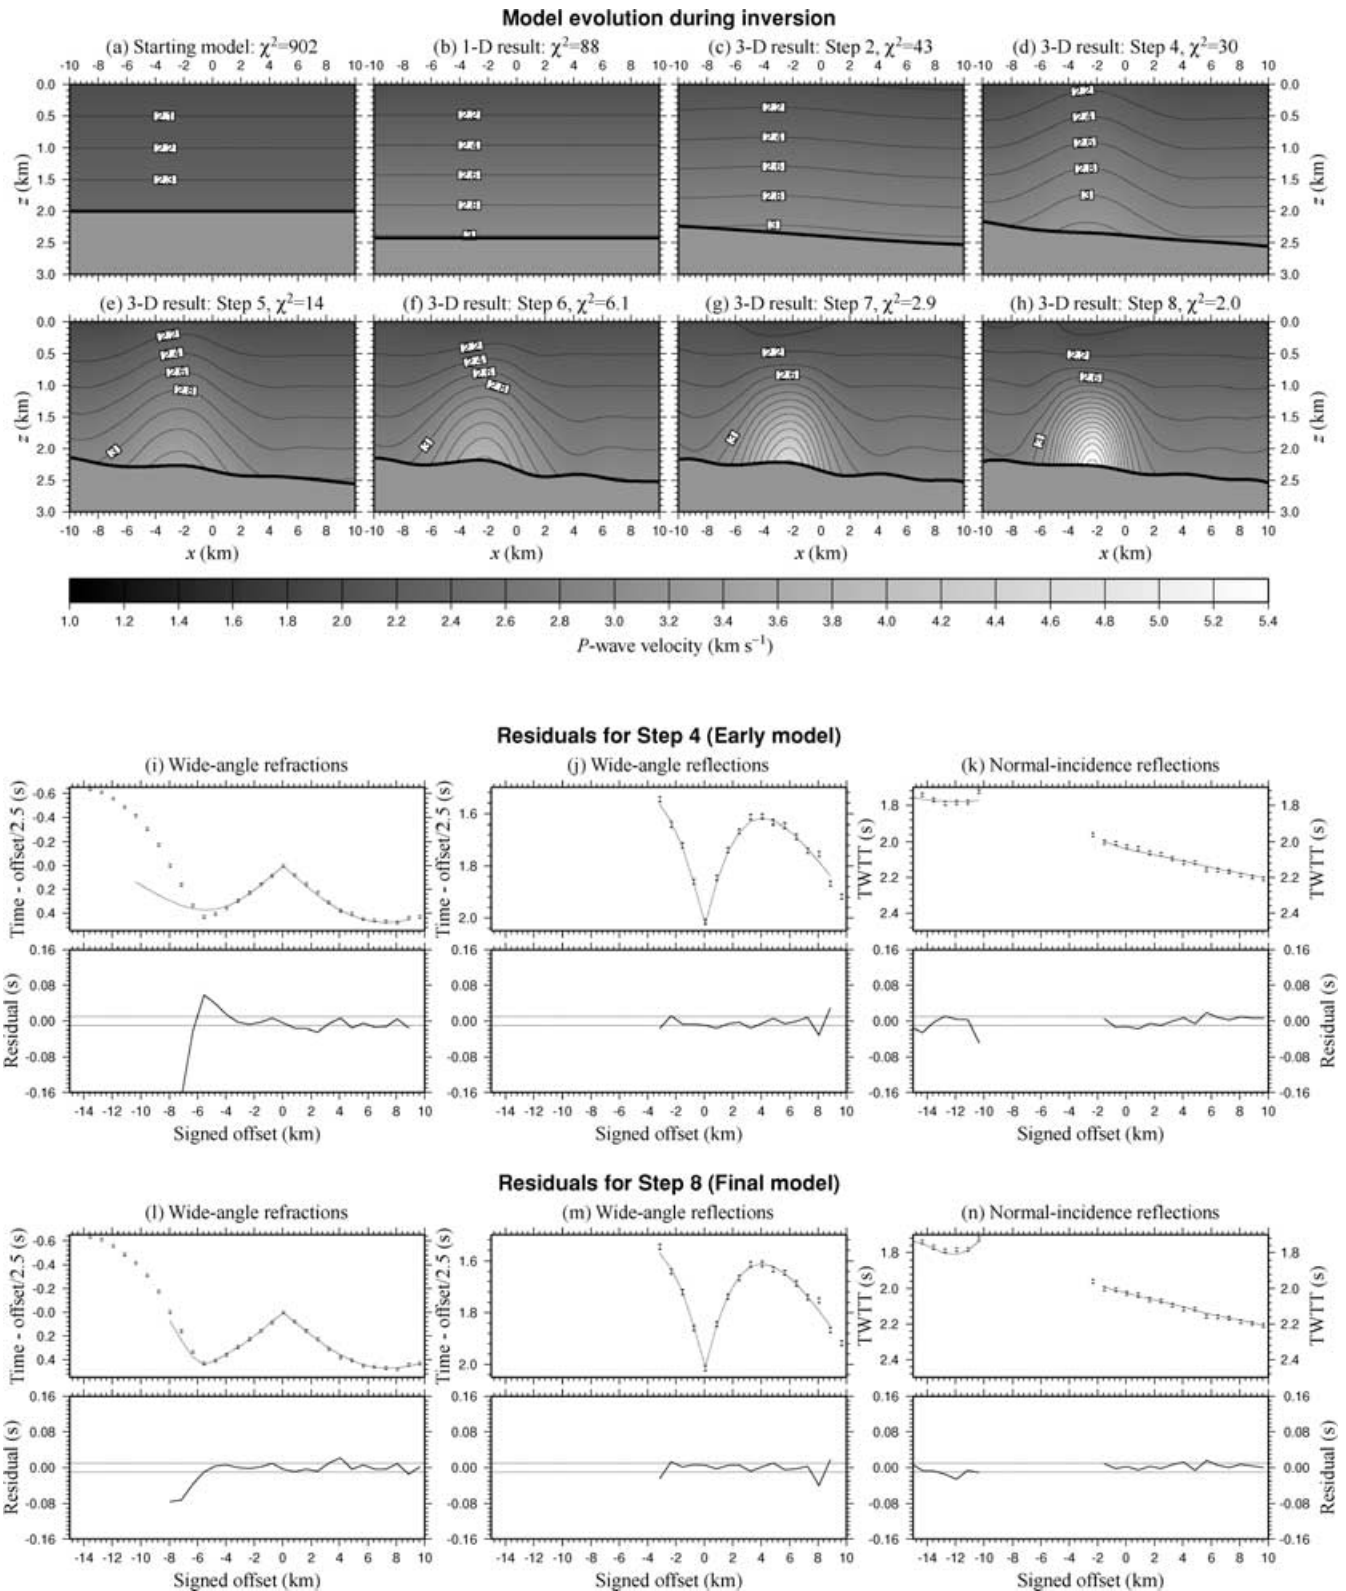


Figure 5. Evolution of the salt dome inversion showing (a) the starting model, (b) the results from the 1-D optimization and (c)–(h) results from the 3-D optimization at successively weaker regularization; also residuals from the salt dome inversion for (i,l) wide-angle refracted, (j,m) wide-angle reflected and (k,n) normal-incidence arrivals, showing the traveltime data with error bars (points) overlain by synthetic data from step 4 and step 8 (lines), and the residuals plotted against the uncertainties in the data.

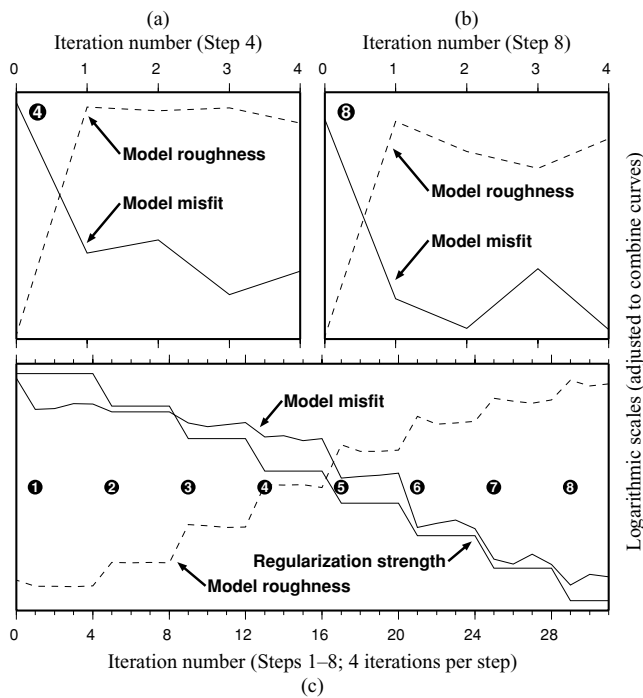


Figure 6. Convergence statistics for synthetic inversion: Plots (a) and (b) illustrate the change in model misfit and model roughness over the four iterations of steps 4 and 8 of the inversion, showing a dramatic change in misfit and roughness during the first iteration followed by smaller adjustments characterized by a slight re-smoothing of the model and small oscillating changes in both measures caused by the trade-off between roughness and misfit during inversion. Plot (c) shows how the same quantities change throughout the eight inversion steps as the regularization strength steadily decreases.

4 APPLICATION TO REAL DATA

The inversion method was applied to a real 3-D seismic data set acquired over the accretionary prism on the Cascadia Margin at approximately 49°N. This is a region in which a strong and continuous bottom-simulating reflector (BSR), the reflector normally associated with the presence of gas hydrate, has been observed in seismic reflection profiles (e.g. Hyndman *et al.* 1994).

The 3-D seismic data set was acquired in June 1993 on the CSS John P. Tully as a collaborative experiment between the University of Victoria, British Columbia and the University of Cambridge. Its primary objective was to image and to obtain constraints on the seismic velocity structure above and immediately below the BSR. The survey was centred around Ocean Drilling Program (ODP) Site 889B (Westbrook *et al.* 1994). The data set used in this analysis was the result of two consecutive deployments of five ocean-bottom hydrophones (OBHs) on which wide-angle data were recorded on a number of profiles for offsets up to 13 km with simultaneous acquisition of single-channel reflection data from a hydrophone at the surface. The second deployment provided constraints at a higher resolution than the first (from a closer line and OBH spacing) with a shorter maximum offset. A single 120 cu. in. airgun source was used, for which the dominant source frequencies were 60–90 Hz. The most serious side-effect of this was the introduction of a strong bubble pulse throughout the entire data set. The hydrophone used to collect the single-channel data recorded at a sampling interval of 2 ms and the OBHs recorded at a sampling interval of 4 ms.

A subset of these data, recorded along a single acquisition line, was modelled by Hobro *et al.* (1998) in two dimensions using the

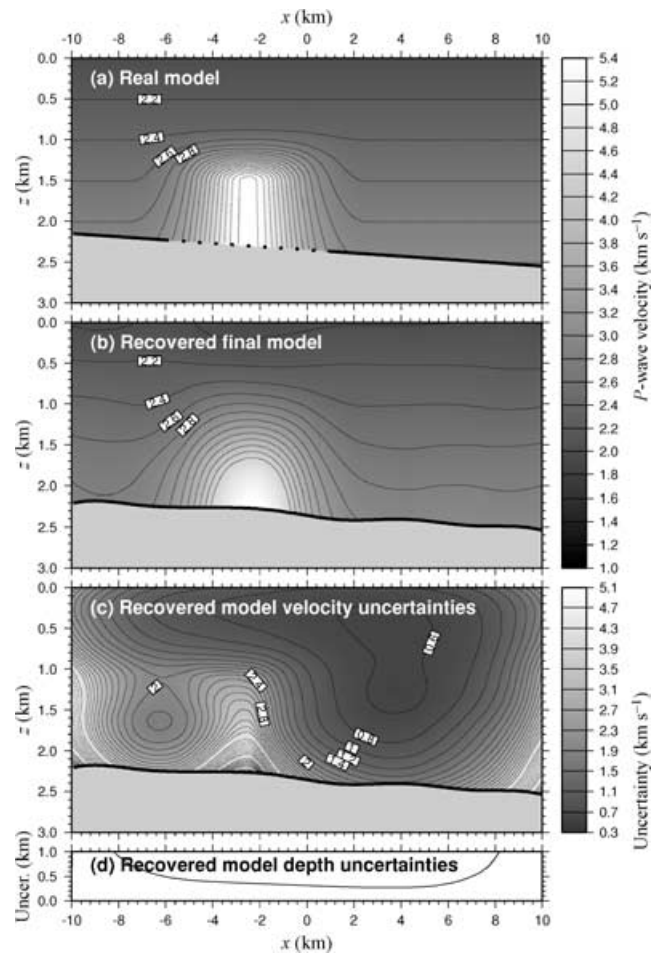


Figure 7. Vertical sections (at $y = -2.5$) through (a) the real salt dome model, (b) the final model from the inversion and (c) and (d) uncertainties for the final model.




Figure 8. Estimated uncertainty (solid line) and actual discrepancy (dotted line) for the final model plotted in Fig. 7, evaluated at $y = -2.5$ and $z = -1.5$.

method of McCaughey & Singh (1997). In this study, a 2-D velocity model of a highly-deformed region of the survey area was successfully obtained from a combination of reflection and refraction seismic data.

The geometry of the two data sets analysed here is presented in Fig. 11(a). In Deployment A, 8 lines were shot with maximum source–receiver offsets of 10–13 km. Deployment B contained 16 lines with maximum offsets of 6–8 km. In both deployments, the instrument positions and line geometry were designed to provide good 3-D velocity constraints on a volume of shallow sediment that




Figure 9. An illustration of the ringing effect produced by regularized least-squares when modelling well-constrained steep velocity gradients. A seismic source is located at the surface and direct arrival traveltimes are obtained for a series of receivers spaced regularly at 10 m depth intervals vertically beneath the source for the model plotted in (a). These times are then inverted using the regularized least-squares method at decreasing regularization strength (b–h). The velocity above the discontinuity at 0.5 km is well constrained, but below the discontinuity, interdependence in the model results in ringing. The asymmetry in this models obtained is due to the asymmetry in the geometry of the experiment.

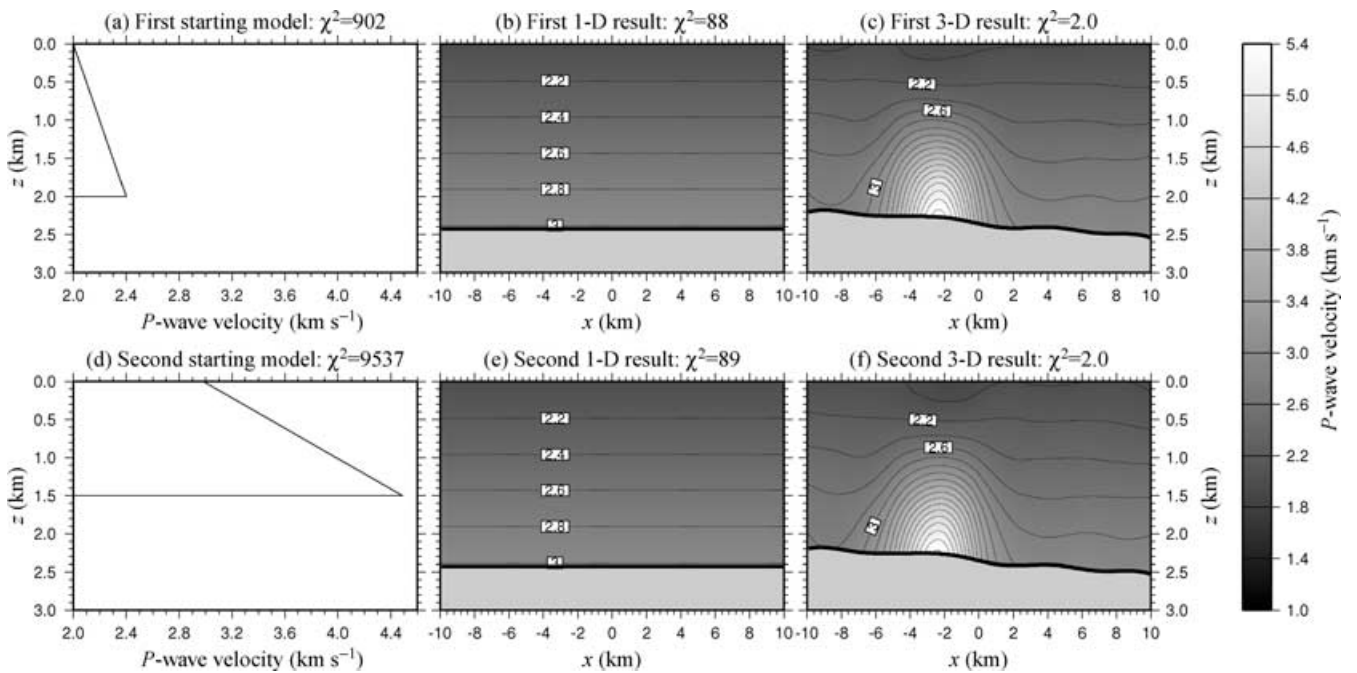


Figure 10. Salt dome inversion results from two different starting models (a) and (d). During the 1-D optimization, the models converge to two indistinguishable solutions (b) and (e). The final results (c) and (f) are almost identical.

was centred on Site 889B. The shot spacing, which nominally remained constant throughout the survey, was approximately 35 m. A total of 112 of the 120 potential wide-angle OBH sections were recorded successfully and 99 of these yielded data that provided useful reflected and/or refracted traveltimes. Samples of the wide-angle and single-channel seismic sections are plotted in Figs 11(b) and (c).

Traveltimes were picked from the OBH data (from horizons interpreted as BSR reflections and energy refracted in the sediments above the BSR) and the single-channel reflection data (seabed and BSR reflections). The inversion method described here was used initially to construct an accurate model of the seabed using seabed picks from all the single-channel data (along all the lines in Fig. 11a) and a water velocity model based on conductivity–temperature–depth profiles. A seabed model more accurate than models provided by bathymetry data was required because the region contained large and complex variations in seabed topography (up to 1 km in depth).

In contrast, the target of the modelling process—the region between the seabed and the BSR—was only 200 m thick.

The seabed was parametrized as an interface spanning 25 km \times 29 km, mapped using 100 \times 100 depth nodes. These 10 000 model parameters were constrained in an inversion for the seabed structure by 9952 seabed reflection picks, obtained from the single-channel data in both deployments. A minimum-structure model that provided a satisfactory fit to all the data was obtained. This seabed model illustrates well the properties of the regularized inversion approach; areas of the model that are sparsely sampled by the data are interpolated smoothly, whereas complex structures exist where they are required to explain the data.

A sub-surface sedimentary velocity layer was defined over a 10 km \times 11 km \times 600 m volume that covered the well-sampled region of the model. This volume was parametrized using 30 \times 30 \times 12 velocity nodes. The layer was bounded at its base by an interface that represented the BSR, which spanned the same lateral region as




Figure 11. (a) A plot of acquisition lines followed during the two wide-angle surveys with the positions of the OBHs indicated by dots; (b) sample OBH section reduced by ‘flattening’ the direct arrival, which shows the main features visible throughout the wide-angle data; (c) sample single-channel seismic section showing the seabed and BSR reflections.

the layer, and was parametrized by 30×30 depth nodes. The node density was chosen to reflect the maximum expected degree of structural variation in the BSR and the seismic velocity field above. The expected complexity was estimated by examining the reflection data and considering the sampling limitations (e.g. Fresnel zone widths). Had a good fit to the data not been achieved using this parameter density, the data would have been re-modelled using a higher depth and velocity node density.

This model was used in a simultaneous inversion for sub-surface velocity and the depth of the BSR, with the water layer and seabed model fixed. A total of 12 446 traveltimes were used (3042 single-channel BSR reflection picks, 3866 picks from turning rays above the BSR in the OBH data and 5538 BSR reflection picks from the OBH data). Since the general data quality and the clarity and coherence of individual horizons varied significantly across the data set, appropriate uncertainties were assigned individually to each data phase in each seismic section. Most of the single-channel reflection picks were assigned an uncertainty of ± 2.5 ms, an estimate of the picking uncertainty due to the bandwidth of the data recorded. In regions where the BSR was not strong, and in which the horizon, although identifiable, could not be picked without intro-

ducing a small amount of noise into the picked data, this uncertainty was increased to ± 4 ms. The normal-incidence constraints on the BSR constitute a dense coverage throughout the focal region of the survey (from $x \approx 11$ km to 17 km and from $y \approx 17$ km to 20 km), and a more sparse coverage elsewhere. Uncertainties for the BSR reflection phase in the OBH data ranged from ± 4 ms (the minimum uncertainty constrained by the bandwidth of the data) to ± 12 ms (a large uncertainty assigned where the precise location of the reflection was not unambiguously identifiable). For the later refracted arrival in the OBH data, uncertainties ranged from ± 5 ms to ± 8 ms, and for the first arrival (which was in general less well-defined than the two later phases), from ± 12 ms to ± 20 ms. This data set provided a demanding test for a tomographic inversion method since the high velocity gradients and the presence of a strong shallow reflector generated a complex set of refracted and reflected horizons.

In order to minimize any errors introduced into the results due to incorrect positioning of the OBHs, all instrument positions were projected vertically to a distance 3 m above the seabed (the correct height for the instruments used in this experiment). The traveltimes picked from the reduced wide-angle data sections (picked relative to the water wave arrival) were then transformed to absolute

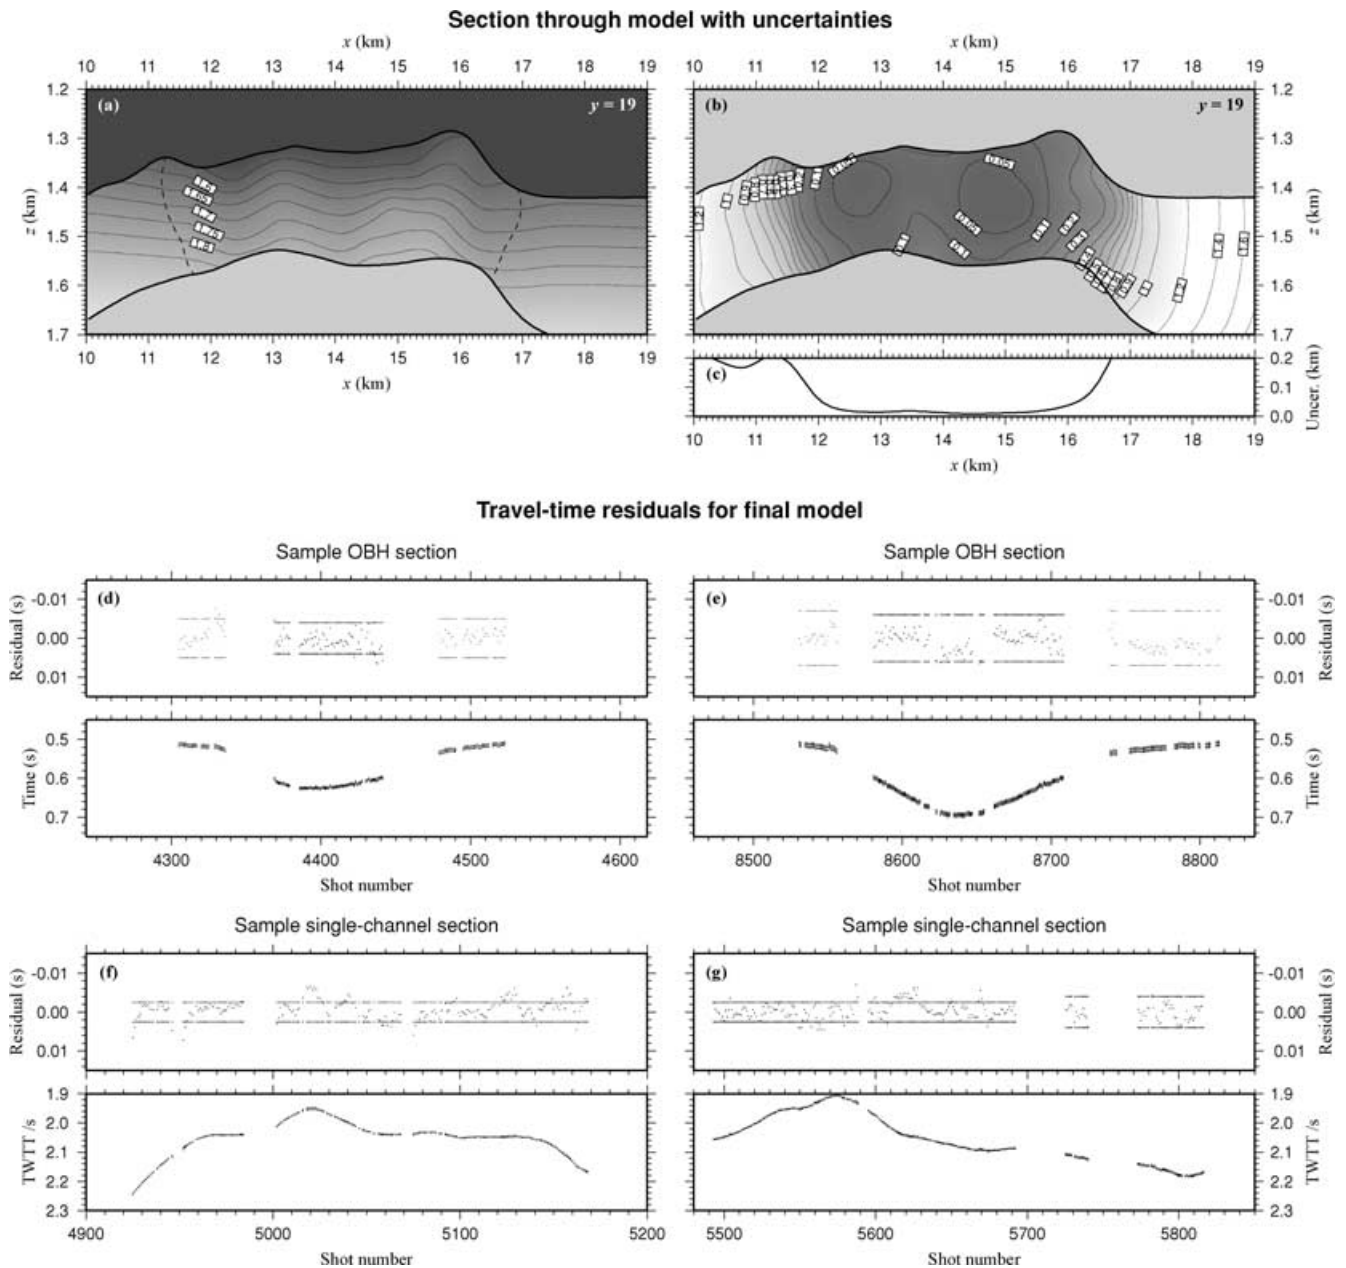


Figure 12. (a) a section through the final model showing variations in the velocity structure above the BSR (the dashed line marks the region in which velocities are well-constrained) with (b) velocity and (c) depth uncertainties; (d), (e), (f) and (g) traveltime residuals from sections of wide-angle and single-channel data, illustrating the quality of fit. Real and synthetic data are superimposed, and above each plot the residuals (black and grey dots) are compared with the uncertainties in each phase (refractions in grey and reflections in black).

traveltimes by adding a direct-arrival traveltimes obtained by ray-tracing through the water column between the appropriate shot-instrument pair.

The inversion method was applied to these data, beginning with a 1-D inversion to determine the average BSR depth, sedimentary P -velocity and velocity gradient, then allowing the regularization strength to decrease step by step as illustrated for the synthetic inversion in Fig. 6. The large variations in topography across the model caused significant lateral velocity gradients to emerge as the 1-D result converged to a smooth 3-D intermediate model. A stable model convergence was achieved and a final model was chosen (after six steps) at the point where systematic components had been eliminated from the vast majority of data residuals (Figs 12d, e, f and g), leaving

only random noise. This indicated that the point had been reached where almost all structure in the data had been modelled without modelling any noise. The data residuals for the final model produced an overall fit of $\chi^2 = 0.76$. A vertical cross-section through the most complex region of the final model is plotted in Fig. 12(a), with the corresponding plot of uncertainties in Figs 12(b) and (c) derived using the method of Section 2.4. This model cross-section illustrates how the inversion method has resolved detailed structure in the well-constrained central region of the model, whereas the smoothing regularization has stabilized the poorly-constrained model edges. The region of well-constrained seismic velocities marked in Fig. 12(a) is obtained from an isoline in the uncertainty function of Fig. 11(e) (at an uncertainty of $\pm 0.6 \text{ km s}^{-1}$).

A detailed analysis of this data set is presented by Hobro (1999). Detailed descriptions of the data set, its analysis, and interpretation of the final model are not presented here—the analysis of this data set is included in this paper in order to illustrate the ability of the inversion method described to model a detailed 3-D structure from real seismic data. The inversion method and data set are well-suited since, in this highly-deformed region with steep velocity gradients, the combination of reflection and refraction data provide considerably better constraints on the seismic velocity structure than either individual data type. In this example, a single iteration typically took around 4 minutes to complete the forward modelling phase, and a similar time to complete the inversion phase on a single Intel Pentium II 300 MHz CPU.

Further applications to real data sets have been presented by Funck *et al.* (2000) and Tong *et al.* (2002).

5 CONCLUSION

A tomographic inversion method has been presented that obtains 3-D velocity–depth models from a range of reflection and refraction seismic traveltime data types. The method makes simultaneous use of reflection and refraction data in a constrained inversion that requires minimal user interaction and selects smooth minimum-structure models. The use of smooth basis functions in the model parametrization reduces the non-linearity in the forward-modelling and inversion schemes and consequently increases the robustness and efficiency of the method. Ray perturbation theory provides a semi-analytical method for forward-modelling through these smooth media. An inversion approach is adopted that allows the bulk properties of the model (average velocities, velocity gradients and average interface depths) to be modelled initially with more detailed structure emerging as the inversion process continues. This provides an effective and objective approach to solving the highly non-linear tomographic problem whilst maintaining a high degree of independence from the starting model.

A series of tests on synthetic data were presented, from which it was concluded that:

- (i) the method recovered a realistic target structure well from a set of noisy traveltime data,
- (ii) the uncertainty analysis provided a useful assessment of the quality of constraint across the model, and
- (iii) the final model recovered did not depend on the starting model used.

The method was applied to a real 3-D seismic data set containing coincident seismic reflection and wide-angle data, the target of which was a region of the methane hydrate stability zone in the Cascadia Margin. A minimum-structure velocity–depth model of the region was obtained that provided a good fit to the traveltime data from all the phases picked. The results confirmed the ability of this method to model a complex seismic structure successfully at an efficient level of parametrization.

ACKNOWLEDGMENTS

We would like to thank George Spence and all involved in the Tully 1993 cruise for acquiring the data set presented here, which was funded by NSERC in Canada and NERC in the UK. J. Hobro was supported by a NERC research studentship and T. Minshull by a Royal Society University Research Fellowship. We thank Jay Pulliam and Colin Zelt for thorough reviews.

REFERENCES

- Aki, K. & Lee, W.H.K., 1976. Determination of three-dimensional velocity anomalies under a seismic array using first *P* arrival times from local earthquakes, *J. geophys. Res.*, **81**, 4381–4399.
- Buchanan, J.L. & Turner, P.R., 1992. *Numerical Methods and Analysis*, McGraw-Hill, Inc.
- Catchings, R.D. & Mooney, W.D., 1988. Crustal structure of the Columbia Plateau: evidence for continental rifting, *J. geophys. Res.*, **93**, 459–474.
- Červený, V., 1987. Ray tracing algorithms in three-dimensional laterally varying layered structures, in *Seismic Tomography, with Applications in Global Seismology and Exploration Geophysics*, ed. Nolet, G., Reidel, Dordrecht.
- Červený, V., Molotkov, I.A. & Psencik, I., 1977. *Ray Method in Seismology*, University of Karlova Press, Prague.
- Chapman, C.H., 1985. Ray theory and its extensions: WKBJ and Maslov seismograms, *J. Geophys.*, **58**, 27–43.
- Constable, S.C., Parker, R.L. & Constable, C.G., 1987. Occam's inversion: A practical algorithm for generating smooth models from electromagnetic sounding data, *Geophysics*, **52**, 289–300.
- Eberhart-Phillips, D., 1986. Three-dimensional velocity structure in northern California Coast Ranges from inversion of local earthquake arrival times, *Bull. seism. Soc. Am.*, **76**, 1025–1052.
- Firbas, P., 1987. Tomography from seismic profiles, in *Seismic Tomography, with Applications in Global Seismology and Exploration Geophysics*, ed. Nolet, G., Reidel, Dordrecht.
- Funck, T., Loudon, K.E., Wardle, R.J., Hall, J., Hobro, J.W.D., Salisbury, M.H. & Muzzatti, A., 2000. Three-dimensional structure of the Torngat Orogen (NE Canada) from active seismic tomography, *J. geophys. Res.*, **105**, 23 403–23 420.
- Ghose, S., Hamburger, M.W. & Virieux, J., 1998. Three-dimensional velocity structure and earthquake locations beneath the northern Tien Shan of Kyrgyzstan, central Asia, *J. geophys. Res.*, **103**, 2725–2748.
- Henry, W.J., Mechie, J., Maguire, P.K.H., Khan, M.A., Prodehl, C., Keller, G.R. & Patel, J., 1990. A seismic investigation of the Kenya Rift Valley, *Geophys. J. Int.*, **100**, 107–130.
- Hestenes, M. & Stiefel, E., 1952. Methods of conjugate gradients for solving linear systems, *Nat. B. Stand. J. Res.*, **49**, 409–436.
- Hobro, J.W., Minshull, T.A. & Singh, S.C., 1998. Tomographic seismic studies of the methane hydrate stability zone in the Cascadia Margin, in *Gas Hydrates: Relevance to World Margin Stability and Climate Change*, Vol. 137, pp. 133–140, Geological Society, London.
- Hobro, J.W.D., 1999. Three-dimensional tomographic inversion of combined reflection and refraction seismic traveltime data, *PhD thesis*, Department of Earth Sciences, University of Cambridge, Cambridge.
- Holbrook, W.S., Reiter, E.C., Purdy, G.M., Sawyer, D., Stoffa, P.L., Austin, J.A., Oh, J. & Makris, J., 1994. Deep structure of the U.S. Atlantic continental margin, offshore South Carolina, from coincident ocean bottom and multichannel seismic data, *J. geophys. Res.*, **99**, 9155–9178.
- Hole, J.A., 1992. Nonlinear high-resolution three-dimensional seismic travel time tomography, *J. geophys. Res.*, **97**, 6553–6562.
- Huang, H., Spencer, C. & Green, A., 1986. A method for the inversion of refraction and reflection traveltimes for laterally varying velocity structures, *Bull. seism. Soc. Am.*, **76**, 837–846.
- Hyndman, R.D., Spence, G.D., Yuan, T. & Davis, E.E., 1994. Regional geophysics and structural framework of the Vancouver Island margin accretionary prism, in *Proc. ODP, Init. Rept.*, 146, Part 1. 399–419, eds Westbrook, G.K., Carson, B. & Musgrave, R.J., *et al.*
- Korenaga, J., Holbrook, W.S., Kent, G.M., Keleman, P.B., Detrick, R.S., Larsen, H.C., Hopper, J.R. & Dahl-Jensen, T., 2000. Deep structure of the U.S. Atlantic continental margin, offshore South Carolina, from coincident ocean bottom and multichannel seismic data, *J. geophys. Res.*, **105**, 21 591–21 614.
- Lutter, W.J., Nowack, R.L. & Braile, L.W., 1990. Seismic imaging of upper crustal structure using traveltimes from the PASSCAL Ouachita experiment, *J. geophys. Res.*, **95**, 4621–4631.

- McCaughey, M. & Singh, S.C., 1997. Simultaneous velocity and interface tomography of normal-incidence and wide-aperture seismic traveltimes data, *Geophys. J. Int.*, **131**, 87–99.
- Moser, T.J., 1991. Shortest path calculations of seismic rays, *Geophysics*, **56**, 59–67.
- Ni, J.F., Ibenbrahim, A. & Roecker, S.W., 1991. Three-dimensional velocity structure and hypocenters of earthquakes beneath the Hazara arc, Pakistan: Geometry of underthrusting Indian plate, *J. geophys. Res.*, **96**, 19 865–19 877.
- Nolet, G., 1987. Seismic wave propagation and seismic tomography, in *Seismic Tomography, with Applications in Global Seismology and Exploration Geophysics*, ed. Nolet, G., Reidel, Dordrecht.
- Robertsson, J.O.A., Blanch, J.O. & Symes, W.W., 1994. Viscoelastic finite-difference modeling, *Geophysics*, **59**, 1444–1456.
- Shaw, P.R. & Orcutt, J.A., 1985. Waveform inversion of seismic refraction data and applications to young Pacific crust, *Geophys. J. R. astr. Soc.*, **82**, 375–414.
- Spence, G.D., Clowes, R.M. & Ellis, R.M., 1985. Seismic structure across the active subduction zone of western Canada, *J. geophys. Res.*, **90**, 6754–6772.
- Tarantola, A., 1987. *Inverse Problem Theory: methods for data fitting and model parameter estimation*, Elsevier, Amsterdam.
- Tikhonov, A.N. & Arsenin, V.Y., 1977. *Solutions of Ill-posed Problems*, Wiley, New York.
- Tong, C.H. *et al.*, 2002. Asymmetric melt sills and upper crustal construction beneath overlapping ridge segments, *Geology*, **30**, 83–86.
- Toomey, D.R., Solomon, S.C. & Purdy, G.M., 1994. Tomographic imaging of the shallow crustal structure of the east Pacific rise at 9°30' N, *J. geophys. Res.*, **99**, 24 135–24 157.
- Virieux, J. & Farra, V., 1991. Ray tracing in 3-D complex isotropic media: An analysis of the problem, *Geophysics*, **56**, 2057–2069.
- Wang, B. & Braile, L.W., 1996. Simultaneous inversion of reflection and refraction seismic data and application to field data from the northern Rio Grande rift, *Geophys. J. Int.*, **125**, 443–458.
- Westbrook, G.K., Carson, B., Musgrave, R.J. *et al.*, eds., 1994. *Proc. ODP, Init. Rept.*, **146**, (Part 1).
- Williamson, P.R., 1990. Tomographic inversion in reflection seismology, *Geophys. J. Int.*, **100**, 255–274.
- Zelt, C.A. & Barton, P.J., 1998. Three-dimensional seismic refraction tomography: A comparison of two methods applied to data from the Faroe Basin, *J. geophys. Res.*, **103**, 7187–7210.
- Zelt, C.A. & Smith, R.B., 1992. Seismic traveltimes inversion for 2-D crustal velocity structure, *Geophys. J. Int.*, **108**, 16–34.
- Zhang, J., ten Brink, U.S. & Toksöz, M.N., 1998. Nonlinear refraction and reflection traveltimes tomography, *J. geophys. Res.*, **103**, 29 743–29 757.