Exploration Geophysics (2006) ***37****, 322–330*

Seismic wavefront tracking in 3D heterogeneous media: applications with multiple data classes

Nicholas Rawlinson [1] **Marthijn de Kool** [2] **Malcolm Sambridge** [1]

Key Words: fast marching method, tomography, teleseismic, refracted wave, reflected wave, wavefront tracking.

ABSTRACT

We demonstrate the potential of a recently developed grid-based eikonal solver for tracking phases comprising reflection branches, transmission branches, or a combination of these, in 3D heterogeneous layered media. The scheme is based on a multi-stage fast marching approach that reinitialises the wavefront from each interface it encounters as either a reflection or transmission. The use of spherical coordinates allows wavefronts and traveltimes to be computed at local, regional, and semi-global scales. Traveltime datasets for a large variety of seismic experiments can be predicted, including reflection, wide-angle reflection and refraction, local earthquake, and teleseismic.

A series of examples are presented to demonstrate potential applications of the method. These include: (1) tracking active and passive source wavefronts in the presence of a complex subduction zone; (2) earthquake hypocentre relocation in a laterally heterogeneous 3D medium; (3) joint inversion of wide-angle and teleseismic datasets for P-wave velocity structure in the crust and upper mantle. Results from these numerical experiments show that the new scheme is highly flexible, robust and efficient, a combination seldom found in either grid- or ray-based traveltime solvers. The ability to track arrivals for multiple data classes such as wide-angle and teleseismic is of particular importance, given the recent momentum in the seismic imaging community towards combining active and passive source datasets in a single tomographic inversion.

INTRODUCTION

Over the last few decades, a large variety of grid- and ray-based methods have been developed to solve one of the most common and challenging problems in seismology: the prediction of the source-receiver path and traveltime of seismic energy in heterogeneous media. Common ray-based schemes include shooting, bending, and pseudo bending (e.g., Julian and Gubbins, 1977; Pereyra et al., 1980; Um and Thurber, 1987; Sambridge and Kennett, 1990; Rawlinson et al., 2001), and common grid-based schemes include finite-difference solution of the eikonal equation (e.g., Vidale, 1988; Qin et al., 1992; Cao and Greenhalgh, 1994; Hole and Zelt, 1995; Qian and Symes, 2002) and network or shortest-path methods (e.g., Nakanishi and Yamaguchi, 1986; Moser, 1991; Bai and Greenhalgh, 2005).

Although ray tracing schemes have traditionally been the method of choice in many applications, grid-based schemes have evolved rapidly in recent times and now offer an efficient and robust alternative. Their advantages include: (1) computing traveltimes to all points of a velocity medium, including diffracted arrivals in ray shadow zones; (2) stability in the presence of significant velocity heterogeneity; and (3) consistently finding the first-arrival traveltime in continuous media. Ray tracing schemes may fail to converge to the true two-point path even in mildly heterogeneous media, and provide no guarantee as to whether the located ray corresponds to a first or later arrival. On the other hand, provided a two-point path is located, ray tracing can be efficient, and often produces more accurate traveltimes than any grid-based alternative. However, for many realistic problems, grid-based schemes can be very efficient in computing traveltimes to the required accuracy, particularly when there is a large ratio between the number of sources and receivers or vice versa.

A recently developed grid-based eikonal solver, which is both computationally efficient and highly robust (unconditionally stable for the first-order case), is the so-called fast marching method or FMM (Sethian, 1996; Sethian, 1999; Sethian and Popovici, 1999). FMM implicitly tracks the evolution of first-arrival wavefronts by combining a causal narrow-band evolution scheme with an upwind entropy-satisfying finite difference solution of the eikonal equation. The speed and stability of FMM makes it well suited to problems involving large datasets such as reflection migration and 3D tomography, where it has already been successfully applied (e.g., Popovici and Sethian, 2002; Rawlinson et al., 2006).

A limitation of most grid-based traveltime solvers is that they can only track first arrivals in continuous media; however, it is frequently the case that a majority of the seismic energy from an event arrives later in the wavetrain. In layered media, it is possible to adapt grid-based schemes to track later-arriving reflection and refraction phases. For example, wavefronts can be tracked from both source and receiver to an interface, and Fermat's Principle of stationary time applied to locate reflection points (e.g., Podvin and Lecomte, 1991; Riahi and Juhlin, 1994). In another approach, Hole and Zelt (1995) use the 3D scheme of Vidale (1990) to compute the traveltimes of the reflected wavefield by assuming that the impinging wavefront and reflector are sufficiently smooth to validate a local planar approximation, which then allows Snell's Law to be applied explicitly. Moser (1991) uses a network ray-tracing approach to find reflections by requiring that the shortest path visit a specified set of nodes that lie on the interface.

In an attempt to develop a robust and general scheme for tracking phases comprising any number of reflection and refraction branches in layered media, Rawlinson and Sambridge (2004a, 2004b, 2005) formulate a multi-stage FMM scheme in 2D Cartesian coordinates. The multi-stage FMM propagates a wavefront through a layer until all points of a bounding interface are intersected. A reflected wavefront can then be tracked by reinitialising FMM from the interface back into the incident layer; a transmitted wavefront can be tracked by reinitialising FMM into the adjacent layer. Repeated application of this technique allows phases composed of any number of reflection and refraction

[1] *Research School of Earth Sciences*
Australian National University
Canberra ACT 0200
Australia
Phone: (02) 6125 5512
Facsimile: (02) 6257 2737
Email: nick@rses.anu.edu.au

[2] *Geoscience Australia*

Presented at the 18th ASEG Geophysical Conference & Exhibition (AESC2006), July, 2006. Revised manuscript received 18 October, 2006.

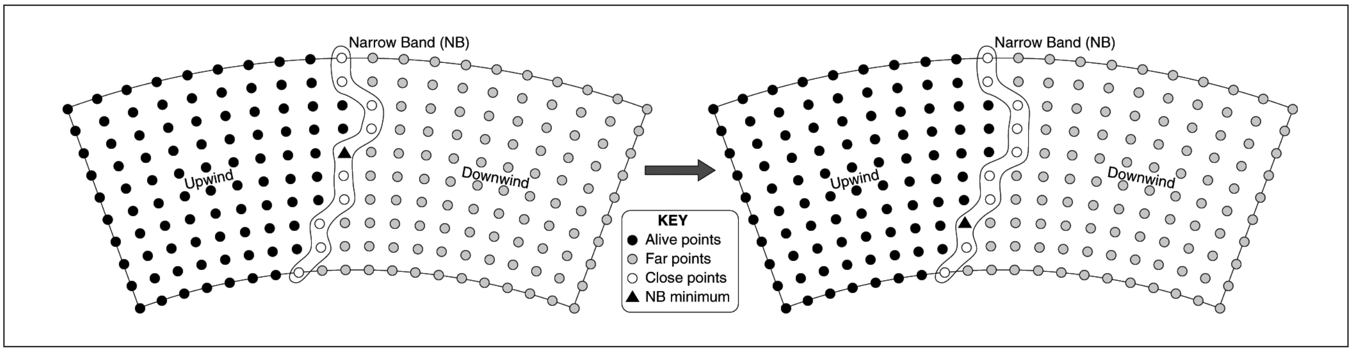

Fig. 1. Principle of the narrow band method for tracking the first arrival wavefront in continuous media. Alive points (upwind of the computational front) have correct traveltimes, Close points have trial traveltimes, and Far points (downwind of the computational front) have no traveltimes computed. The narrow band (NB) evolves by locating the Close point with minimum traveltime, tagging it as Alive, and then computing traveltimes to any adjacent Close or Far points.

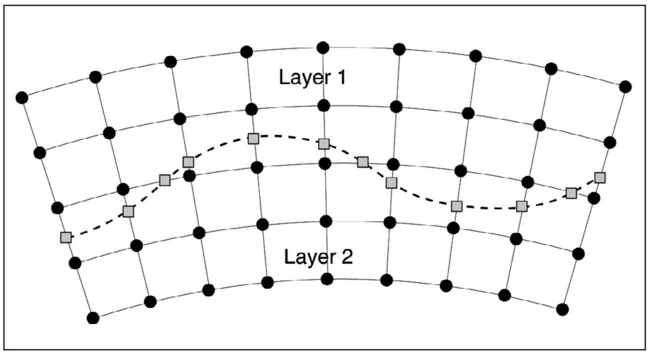

Fig. 2. Schematic diagram (in 2D for clarity) illustrating the relationship between interface nodes (grey squares) and velocity nodes (black dots) in a layered medium. Interface nodes occur at the intersection between the interface surface (thick dashed line) and the mesh lines (thin grey lines) of the regular velocity grid. In 3D, the same principle applies.

branches to be tracked, without the need for explicit application of Fermat's Principle or Snell's Law. Tests of the multi-stage scheme in the presence of highly heterogeneous structure show it to be both computationally efficient and extremely robust.

In a recent paper, de Kool et al. (2006) build on the ideas of Rawlinson and Sambridge (2004a, b) to extend the multi-stage FMM scheme to 3D spherical coordinates. The use of spherical coordinates means that traveltime prediction problems at a variety of scales can be solved. Additional features of the 3D multi-stage scheme include the ability to track multiple reflections from a single interface (e.g., the global PP phase) and P-S conversions at interfaces. A series of numerical tests with a variety of simple and complex velocity models show the new method to be comparable to its 2D counterpart in terms of robustness and computational efficiency.

In this paper, we set out to demonstrate both the flexibility and usefulness of the new method in a variety of applications with synthetic data. The first test demonstrates the ability of the scheme to track reflection, refraction, local earthquake, and teleseismic phases through a complex subduction zone model which features layer pinch-outs, strong interface curvature, and significant lateral velocity heterogeneity. The second test uses an iterative non-linear subspace inversion scheme to relocate earthquake hypocentres in the presence of strong velocity variation. The final test combines refraction, wide-angle reflection, and teleseismic data to perform an inversion for 3D crust, Moho, and upper mantle structure. With this diverse set of examples, we demonstrate that a well-developed grid-based method is more robust and at least as flexible and efficient as any ray-tracing scheme, and is therefore appropriate for a wide class of seismic problems.

THE MULTI-STAGE FAST MARCHING METHOD IN 3D SPHERICAL COORDINATES

A brief description of the 3D multi-stage scheme is given below; for more details refer to Rawlinson and Sambridge (2004a,b) and de Kool et al. (2006). The eikonal equation, which governs the propagation of seismic waves in the high frequency limit, may be written

$$\left|\nabla_{\mathbf{x}} T\right| = s(\mathbf{x}) \ , \tag{1}$$

where ∇_x is the gradient operator, T is traveltime and $s(\mathbf{x})$ is slowness as a function of position **x**. FMM solves equation (1) using upwind entropy satisfying finite differences that naturally deal with wavefront discontinuities that arise from discarding later arriving information. In our case, we use the following scheme, which has been employed by a number of authors, including Sethian and Popovici (1999), Chopp (2001), and Popovici and Sethian (2002):

$$\begin{bmatrix} \max\left(D_a^{-r}T, -D_b^{+r}T, 0\right)^2 \\ \max\left(D_c^{-\theta}T, -D_d^{+\theta}T, 0\right)^2 \\ \max\left(D_e^{-\phi}T, -D_f^{+\phi}T, 0\right)^2 \end{bmatrix}_{i,j,k}^{\frac{1}{2}} = s_{i,j,k} \ , \tag{2}$$

where (i, j, k) are grid increment variables in any orthogonal coordinate system (r,θ,ϕ), and the integer variables a, b, c, d, e, f define the order of accuracy of the upwind finite-difference operator used in each of the six cases. If (r,θ,ϕ) represent a spherical coordinate system, then the first two gradient operators in each of the three directions can be written

$$\begin{aligned} D_1^{-r}T_{i,j,k} &= \frac{T_{i,j,k} - T_{i-1,j,k}}{\delta r} \\ D_2^{-r}T_{i,j,k} &= \frac{3T_{i,j,k} - 4T_{i-1,j,k} + T_{i-2,j,k}}{2\delta r} \\ D_1^{-\theta}T_{i,j,k} &= \frac{T_{i,j,k} - T_{i,j-1,k}}{r\delta\theta} \\ D_2^{-\theta}T_{i,j,k} &= \frac{3T_{i,j,k} - 4T_{i,j-1,k} + T_{i,j-2,k}}{2r\delta\theta} \\ D_1^{-\phi}T_{i,j,k} &= \frac{T_{i,j,k} - T_{i,j,k-1}}{r\cos\theta\delta\phi} \\ D_2^{-\phi}T_{i,j,k} &= \frac{3T_{i,j,k} - 4T_{i,j,k-1} + T_{i,j,k-2}}{2r\cos\theta\delta\phi} \end{aligned} \ . \tag{3}$$

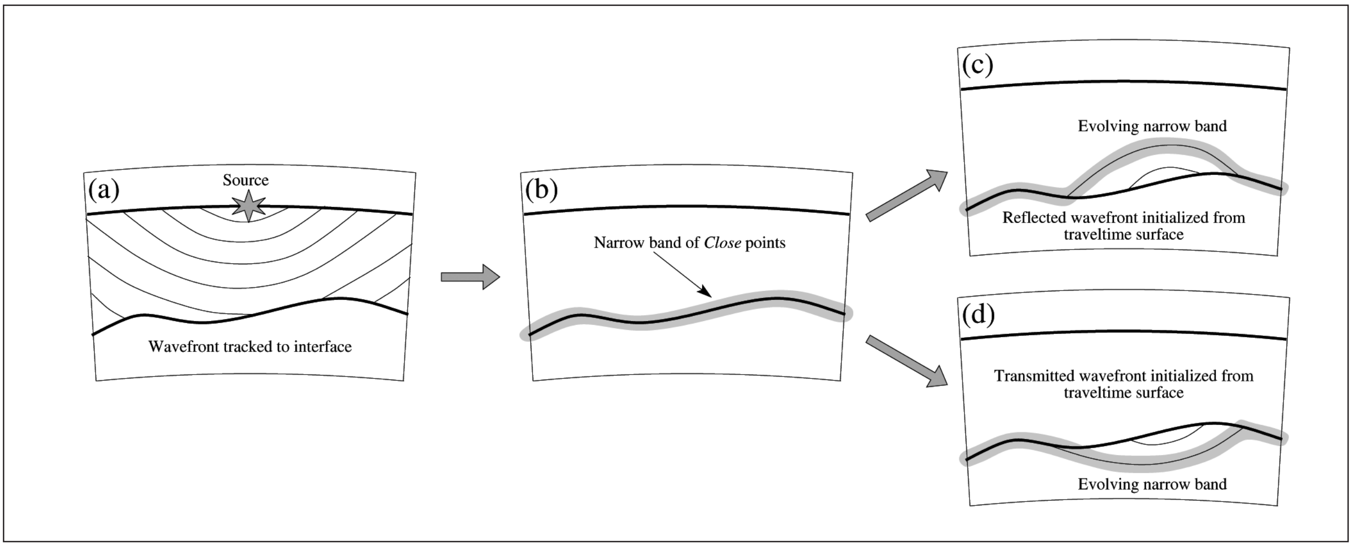

Fig. 3. Principle behind the multi-stage FMM approach. The incident wavefront is tracked to all points on the wavefront, before reinitialisation of FMM in the incident (for reflection) or adjacent (for transmission) layer.

Which operator is used in equation (2) depends on the availability of upwind traveltimes and the maximum order allowed. By default, we use a mixed order scheme which preferentially uses D_2 operators but reverts to D_1 if $T_{i\text{-}2,\,j,\,k}$, $T_{i,\,j\text{-}2,\,k}$, or $T_{i,\,j,\,k\text{-}2}$ is unavailable (e.g., near a point source). Mixed order schemes using higher order operators such as D_3 may also be devised.

Equation (2) describes the finite difference scheme for updating the traveltime associated with a particular grid point, but does not specify the order in which points are to be updated. To satisfy causality, the order should be consistent with the direction of information flow, that is, from smaller to large values of *T*. FMM achieves this by systematically constructing traveltimes in a downwind fashion from known values upwind, using a narrow-band approach (see Figure 1). Close points have trial traveltime values which only become Alive when they are the global minimum value of all Close points. When this occurs, the narrow band is locally evolved to retain the division of the traveltime time field into Alive and Far points. Thus, the shape of the narrow band approximates the shape of the first arrival wavefront, and the idea is to propagate the band through the grid until all points become Alive. The use of a heap-sort algorithm means that FMM has an operation count of O(*N*log*N*) where *N* is the total number of grid points. The fact that the computation cost scales with grid size in this way is responsible for the overall efficiency of FMM.

Rawlinson and Sambridge (2004b) show that in most cases, the dominant error in FMM calculations is accumulated in the source neighbourhood as a result of high wavefront curvature. They show that accuracy can be greatly increased without significant computational cost by introducing a refined grid in the source vicinity, which is resampled to the global coarse grid when the narrow band reaches the refined grid boundary. This scheme is also implemented in the 3D multi-stage version; typical values for the refined grid are an increase in resolution by a factor of 5, and a refined grid extending 50 nodes away from the source in each direction.

The FMM scheme described above is for a regular grid implementation; however, when wave propagation occurs in a layered medium with undulating interfaces, the boundary of each velocity continuum may be irregular. Rawlinson and Sambridge (2004a,b) deal with this problem by using an adaptive triangular meshing scheme in the neighbourhood of each interface to connect interface nodes with adjacent velocity nodes. Interface nodes are defined by the point of intersection between the interface surface and the grid lines of the velocity mesh (see Figure 2). Wavefronts are then propagated through the irregular grid using a first-order version of equation (2) for triangular elements (equivalent to a plane wave approximation). In 3D, an adaptive tetrahedral mesh would be much harder to implement, so instead, updates to interface or surrounding velocity nodes are made by considering impinging plane waves from all possible combinations of neighbouring Alive velocity and interface nodes. The correct wavefront orientation yields the minimum arrival time (see de Kool et al., 2006, for more details).

The principle underlying the multi-stage FMM approach is shown in Figure 3; a wavefront is tracked from the source until all interface nodes become Alive. All Alive points are then reset to Close and FMM is reinitialised into the incident layer for a reflection, or into the adjacent layer for a refraction. This can be repeated any number of times to construct phases comprising any number of refraction or reflection branches. Conversions between P and S waves can be tracked simply by replacing the incident or adjacent velocity field with the appropriate P or S wavespeed model prior to reinitialisation of FMM. Ray paths for any phase type can be found a posteriori by simply integrating along the traveltime field gradient from the receiver back to the source.

A class of phase that cannot be tracked using this multi-stage FMM approach is one that involves consecutive interactions with the same interface (e.g., PP); this is because it no longer represents the interface-intersecting global minimum arrival within the layer. In our implementation of the 3D multi-stage scheme, phases such as PP can be computed by initialising FMM from both the source and receiver and tracking the resulting wavefronts to a common interface. The complete phase can then be obtained by matching the traveltime gradients from the two impinging wavefronts at the interface; where they are equal in magnitude and opposite in sign corresponds to a legitimate reflection point. This approach also allows multiple later-arriving reflections from a single interface to be found if they exist. However, a drawback in finding this class of later arriving phase is that FMM must be initiated from both source and receiver, which increases computation time.

In addition to tracking wavefronts from point sources within the model volume, the multi-stage FMM can also be initialised from a teleseismic wavefront. This is done by computing traveltimes from distant sources to the boundary of the model using, for example,

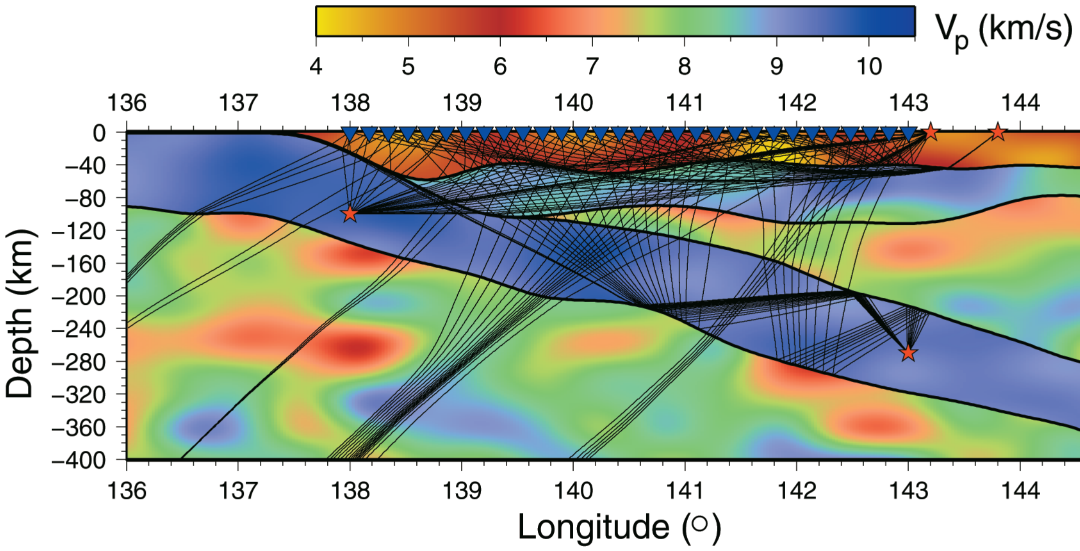

Fig. 4. Teleseismic, local earthquake, reflection, and refraction paths tracked through a complex subduction zone example using the multi-stage FMM scheme. The model volume varies in latitude between 37°S and 44°S, but the model structure only varies with longitude and depth. Red stars denote sources and blue triangles denote receivers.

a global reference model such as ak135 (Kennett et al., 1995). Initialisation of the teleseismic wavefront from the boundary of the model is carried out in the same way as from an interface: the Alive grid points are set to Close and the narrow band evolves from the point with minimum traveltime. Using this approach, any global phase (e.g., P, PcP, PKiKP) can be tracked through the local model. This style of wavefront propagation is often employed in teleseismic tomography (e.g., Graeber et al., 2002; Rawlinson et al., 2006).

EXAMPLES

In the following series of examples, the velocity field within each layer is independently described by a regular grid of nodes in spherical coordinates. These nodes are used as the control vertices of a mosaic of cubic B-spline volume elements, which define the continuum. Layer interfaces are described by a regular mesh in latitude and longitude and use a mosaic of cubic B-spline surface patches to describe the complete interface. Cubic B-spline functions are desirable because they exhibit local control (i.e., changing the value of a single node only effects the interface or velocity field in the vicinity of the node), continuity of the second derivative, and can be rapidly evaluated. This last property is important because in order to use the multi-stage FMM, a propagation grid must be defined over which the narrow band evolves; typically, this will involve many evaluations of the spline function.

Multiple phases in a subduction zone

One of the more complex geological structures encountered at the lithospheric scale is a subduction zone, which involves the edge of one tectonic plate (usually oceanic) subducting beneath another at a convergent margin. The accurate representation of such a structure with a 3D wavespeed model would require that layer pinch-outs, a subducting slab, and strong lateral wavespeed perturbations be included. Although the presence of these complexities makes the task of tracking wavefronts and computing traveltimes a difficult one, it may well be necessary in some circumstances, for instance, in tomographic inversion of multiple datasets. With this in mind, we demonstrate the ability of the multi-stage scheme to track reflection, refraction, local earthquake, and teleseismic phases in a subduction zone setting.

Figure 4 shows a synthetic subduction zone model, comprising an oceanic slab subducting beneath a laterally discontinuous continental mass. Interface topography is relatively complex and lateral variations in wavespeed within a layer are as great as 25%. Although the velocity variations indicated in the scale bar of Figure 4 are probably larger than would be encountered in the Earth, the purpose of this example is to demonstrate the robustness of the scheme. For the sake of clarity, the subduction zone in Figure 4 is 2.5D; that is, it fills a 3D volume but only varies in longitude and depth; this allows ray paths to be meaningfully visualised. Ray paths from five sources – one teleseismic, two local earthquakes within the slab, one refraction, and one reflection – to 30 receivers in-line on the surface, are superimposed on the velocity model in Figure 4. Ray paths from the teleseismic source clearly reflect the distortion experienced by the wavefront as it propagates through the subduction zone. Of the two local earthquake sources, a direct transmission is tracked from the shallower event, while a reflection multiple is tracked from the deeper event. Rays from the easternmost surface source show that even near-critical refraction paths can be retrieved. The trade-off between accuracy and computation time is difficult to quantify in the absence of analytic solutions, but convergence tests involving the progressive increase in density of the propagation grid suggest that all 150 traveltimes can be computed with an RMS error of less than 100 ms in approximately 150 s on a 1.6 GHz Opteron workstation.

Figure 5 shows a second subduction zone example in which variations in structure are fully 3D. In this case, paths for both teleseismic and local earthquake sources are computed. P-S converted phases are shown for the two sources that lie in the subducting slab. The main aim of these subduction zone examples is to highlight both the flexibility and robustness of the multi-stage FMM in 3D. That the stability of the scheme is uninhibited by the complexity of the structure, and that many classes of dataset can be synthesised, means that the multi-stage FMM has the potential to form the basis of a flexible nonlinear tomographic imaging scheme.

Earthquake relocation

The accurate determination of earthquake hypocentres is important for many reasons, including seismic hazard analysis,

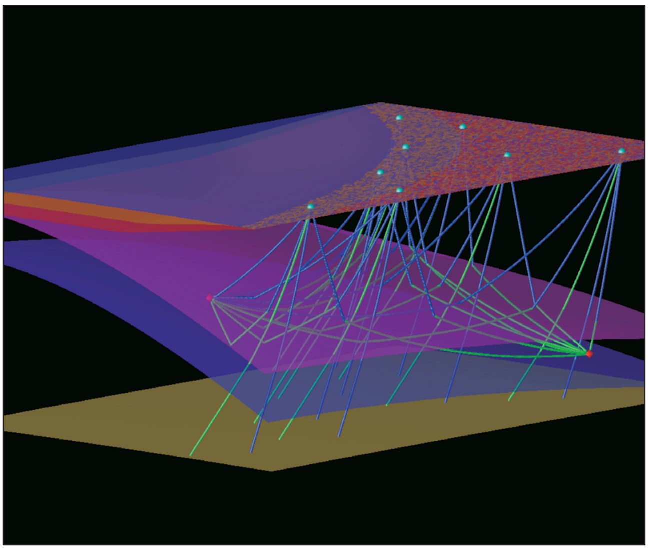

Fig. 5. Tracking converted phases (P-S) through a fully 3D subduction zone from local and teleseismic sources. Green and purple rays denote the two point paths taken by P and S phases respectively.

seismic imaging (e.g., regional tomography), and structural inference (e.g., fault delineation). In this example (Figure 6), we apply a gradient-based iterative non-linear inversion routine, which exploits traveltime derivatives with respect to source location parameters, to relocate (in space and time) a cluster of five local earthquakes in the presence of significant 3D heterogeneity. Figure 6a shows the true location of the earthquakes and Figure 6b shows the initial locations used in the inversion. The synthetic dataset comprises first-arrival traveltimes from all sources to 100 receivers located at the surface. The corresponding raypaths for the true locations are shown in Figure 6a. A source origin time error of 0.25 s is added to all traveltimes in the synthetic dataset. The inversion routine involves iterative application of the multi-stage FMM and a subspace inversion scheme (Kennett et al., 1988; Rawlinson and Sambridge, 2003), which naturally deals with multiple parameter classes. In the subspace scheme, minimisation is carried out simultaneously along several search directions that together span a subspace of the model space. The search directions used in this case are based on the gradient vector in model space and its repeated pre-multiplication by the model space Hessian (see Rawlinson and Sambridge, 2003). For this example, a 20D subspace scheme offers a good compromise between computing time per iteration and rate of convergence.

Figure 6c shows the result after 6 iterations of the inversion scheme; in all five cases, the earthquake hypocentres have been relocated very accurately. Table 1 summarises the spatial and temporal errors associated with the initial and recovered locations. The trade-off between interface depth and source time error is not fully resolved for sources 4 and 5, but this is not surprising given the ray path geometry (Figure 6a). The success of this example in a strongly heterogeneous 3D medium shows that the multi-stage FMM has the potential to be used in routine earthquake relocation algorithms where there is sufficient a priori knowledge of lateral variations in structure. Although an iterative non-linear approach to relocation is employed here, the multi-stage FMM is also particularly well suited to fully non-linear relocation algorithms (e.g., Kennett, 2004) that rely on relatively dense sampling of parameter space. This is because grid-based methods compute traveltimes from a source to all points that define the model volume. Thus, the multi-stage FMM need only be solved once for each receiver, with the resulting traveltime fields stored in look-up tables. Using the principal of traveltime reciprocity, source-receiver traveltimes at an arbitrary point can then be rapidly extracted. In complex velocity models with poor initial locations, a fully non-linear or global search algorithm is preferable to an iterative non-linear one, which is liable to either diverge or become trapped in a local minimum.

Joint inversion of wide-angle and teleseismic traveltimes

The traditional approach to local and regional scale body wave tomography in 3D has been to invert traveltime data of a single class, such as teleseismic (Aki et al., 1977; Humphreys and Clayton, 1990; Graeber et al., 2002; Rawlinson et al., 2006), local earthquake (Eberhart-Phillips, 1986; Walck, 1988; Graeber and Asch, 1999), or wide-angle (Hole et al., 1992; Zelt and Barton, 1998; Rawlinson et al., 2001). Although several efforts have been made in the past to combine multiple datasets in a single inversion (e.g., Ankeny et al., 1986; Sato et al., 1996; Parsons and Zoback, 1997), it is still far from a routine practice. However, with the rapidly increasing volume of seismic data recorded in various regions of the Earth, it is becoming clear that the opportunities for combining multiple datasets are on the rise. The obvious benefit of more than one dataset is the increased path coverage, which is likely to yield more detailed and robust models. For example, wide-angle surveys often provide good coverage of the crust, but sample the upper mantle poorly. On the other hand, teleseismic surveys often result in good path coverage through the upper mantle, but sample the crust poorly.

In the next application of the multi-stage FMM (Figure 7), synthetic wide-angle and teleseismic datasets are simultaneously inverted for the geometry of the Moho and the P wavespeed structure of the crust and upper mantle. The tomographic inversion routine that is used is similar to the one previously applied for earthquake hypocentre relocation, except that interface depth and velocity node parameters are inverted for rather than source parameters. The iterative application of the multi-stage FMM and the subspace inversion scheme accounts for the non-linear relationship between traveltime and wavespeed/interface depth

	Initial source error				Relocated source error			
Source	Latitude (°)	Longitude (°)	Depth (km)	Time (s)	Latitude (°)	Longitude (°)	Depth (km)	Time (s)
1	–0.200	–0.200	10.0	0.250	–0.001	0.000	–0.1	–0.006
2	0.100	0.200	0.0	0.250	0.000	0.001	0.0	0.002
3	–0.100	0.000	–10.0	0.250	0.000	–0.001	0.4	0.011
4	0.200	–0.300	10.0	0.250	–0.001	–0.001	–0.5	–0.023
5	–0.200	0.300	10.0	0.250	0.001	0.002	–0.5	–0.085

Table 1. Summary of results from the earthquake hypocentre relocation example. The initial, recovered, and true source locations are shown in Figure 6.

perturbations. As in the previous example, a 20D subspace scheme is used. In addition to minimising the difference between observed and predicted traveltimes, the inversion also includes damping regularisation, which penalises solution models that are strongly perturbed from the initial or starting model.

Figure 7a shows a plot of the "true" model used to generate the synthetic dataset. It comprises a crustal layer and an upper mantle layer separated by an undulating Moho. Lateral perturbations in wavespeed from a 1D reference model are shown; in this case, the reference model is described by ak135 (Kennett et al., 1995) in the upper mantle, and a constant velocity gradient of 0.016 s^{-1} in the crust, with a velocity of 5.4 km/s at the surface. The synthetic velocity structure comprises a 3D checkerboard pattern in both layers, with a finer scale checkerboard used in the crust, where more detail can be resolved. A total of 14 540 nodes describe the

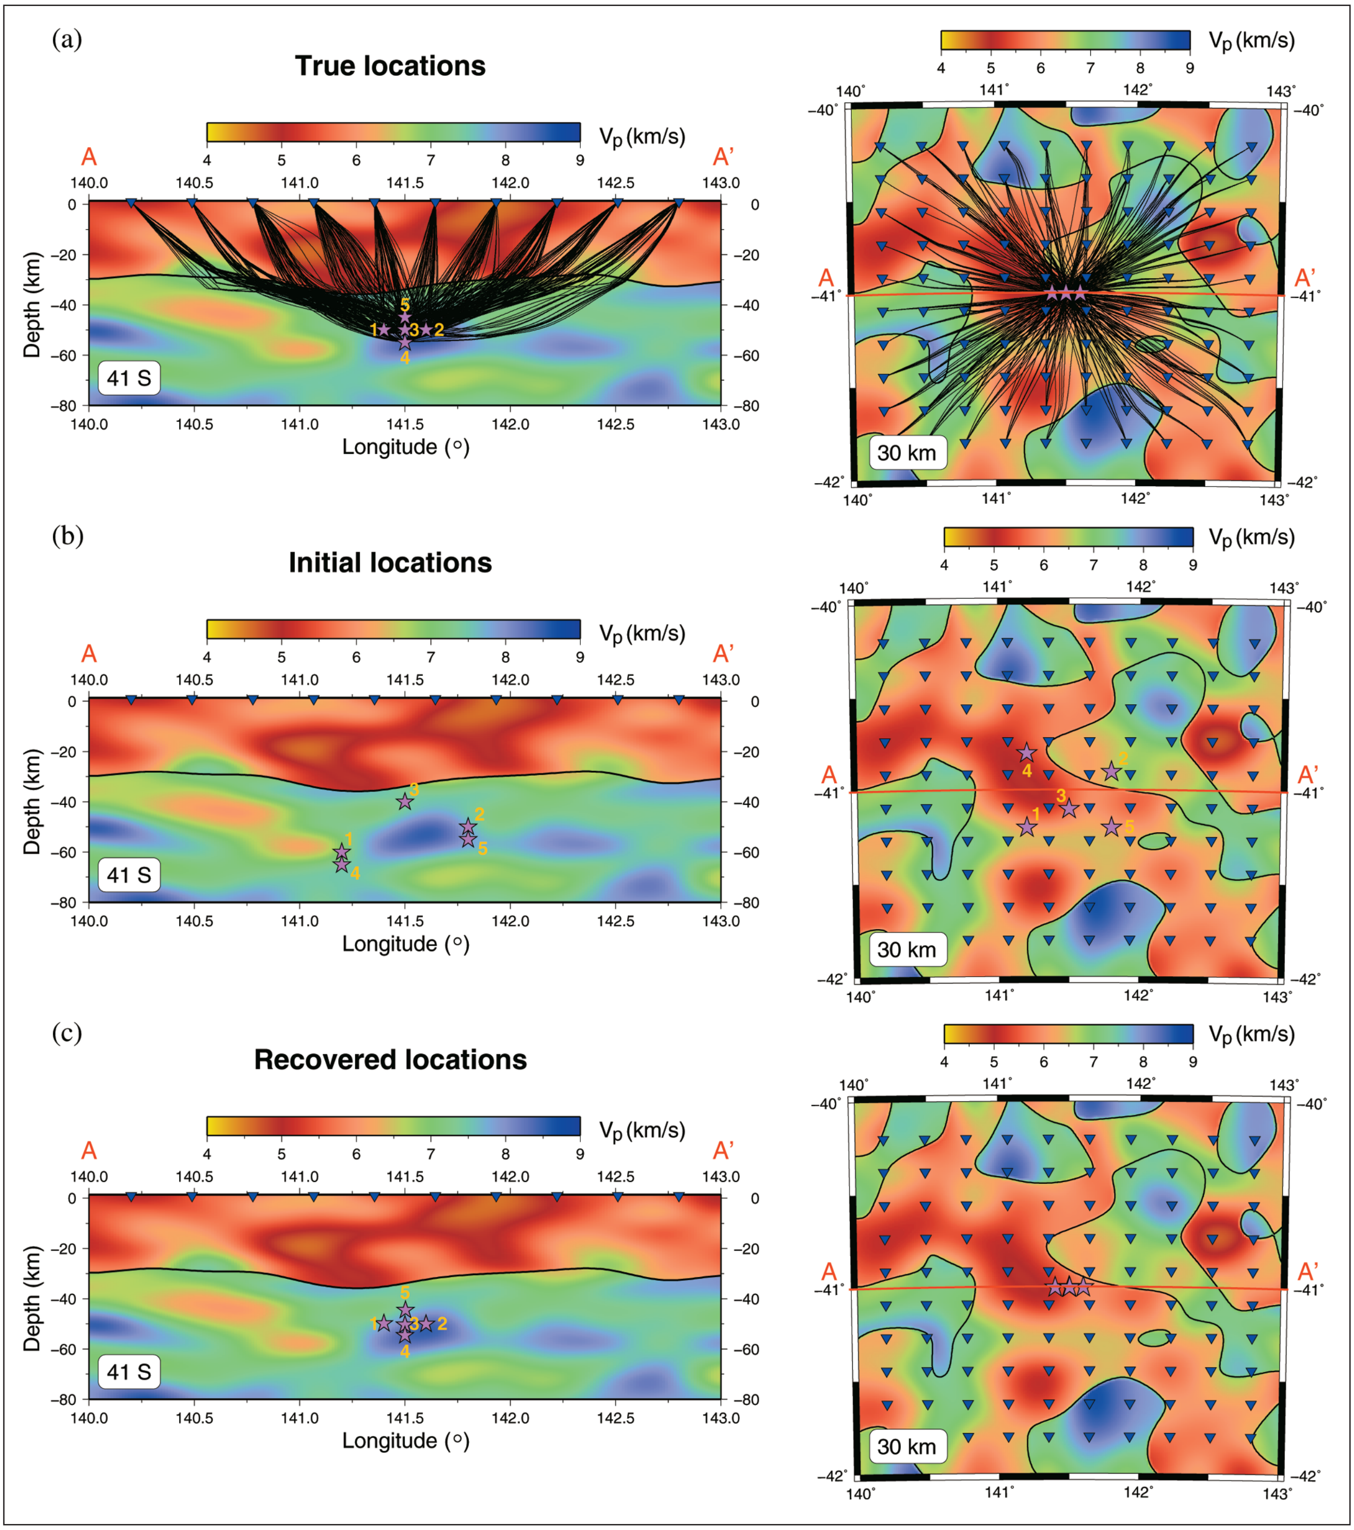

Fig. 6. Results from the earthquake hypocentre relocation example (see also Table 1). (a) True, (b) initial, and (c) recovered locations of five local earthquake sources. Ray paths are also projected on the E-W cross-section and the horizontal section for the true source locations. Note that the horizontal slices at 30 km depth feature velocity discontinuities as a result of the undulating Moho. Magenta stars denote sources and blue triangles denote receivers. Source numbering is consistent with Table 1.

entire 3D velocity field. The synthetic interface structure varies in depth between 33 and 47 km, and is described by 196 depth nodes. Figure 7b shows the laterally invariant starting model used for the tomographic inversion and the associated ray path coverage. The starting velocity model is simply the 1D reference model, and the initial depth of the interface is 40 km. The teleseismic dataset comprises traveltime residuals from 12 distant events recorded by 100 receivers evenly distributed on a grid at the surface. The wide-angle dataset consists of crustal refraction (Pg) and reflection (PmP) traveltimes generated by nine surface sources and recorded by the same 100 receivers. In total, the synthetic dataset comprises 3000 traveltimes.

The tomographic solution model, obtained after six iterations of the non-linear scheme, is shown in Figure 7c. The checkerboard pattern in the crust is clearly recovered, with minimal smearing. The checkerboard pattern in the upper mantle is also well recovered, although there is clear evidence of vertical smearing which prevents the uppermost part of the checkerboard in the bottom layer from being recovered. This behaviour is typical of teleseismic datasets. The inclusion of significantly more teleseismic events, or refraction paths from the surface sources which penetrate the uppermost mantle (Pn), would probably improve the imaging result in this region. The structure of the Moho is generally well recovered, and shows that the trade-off between interface depth and velocity perturbation has been

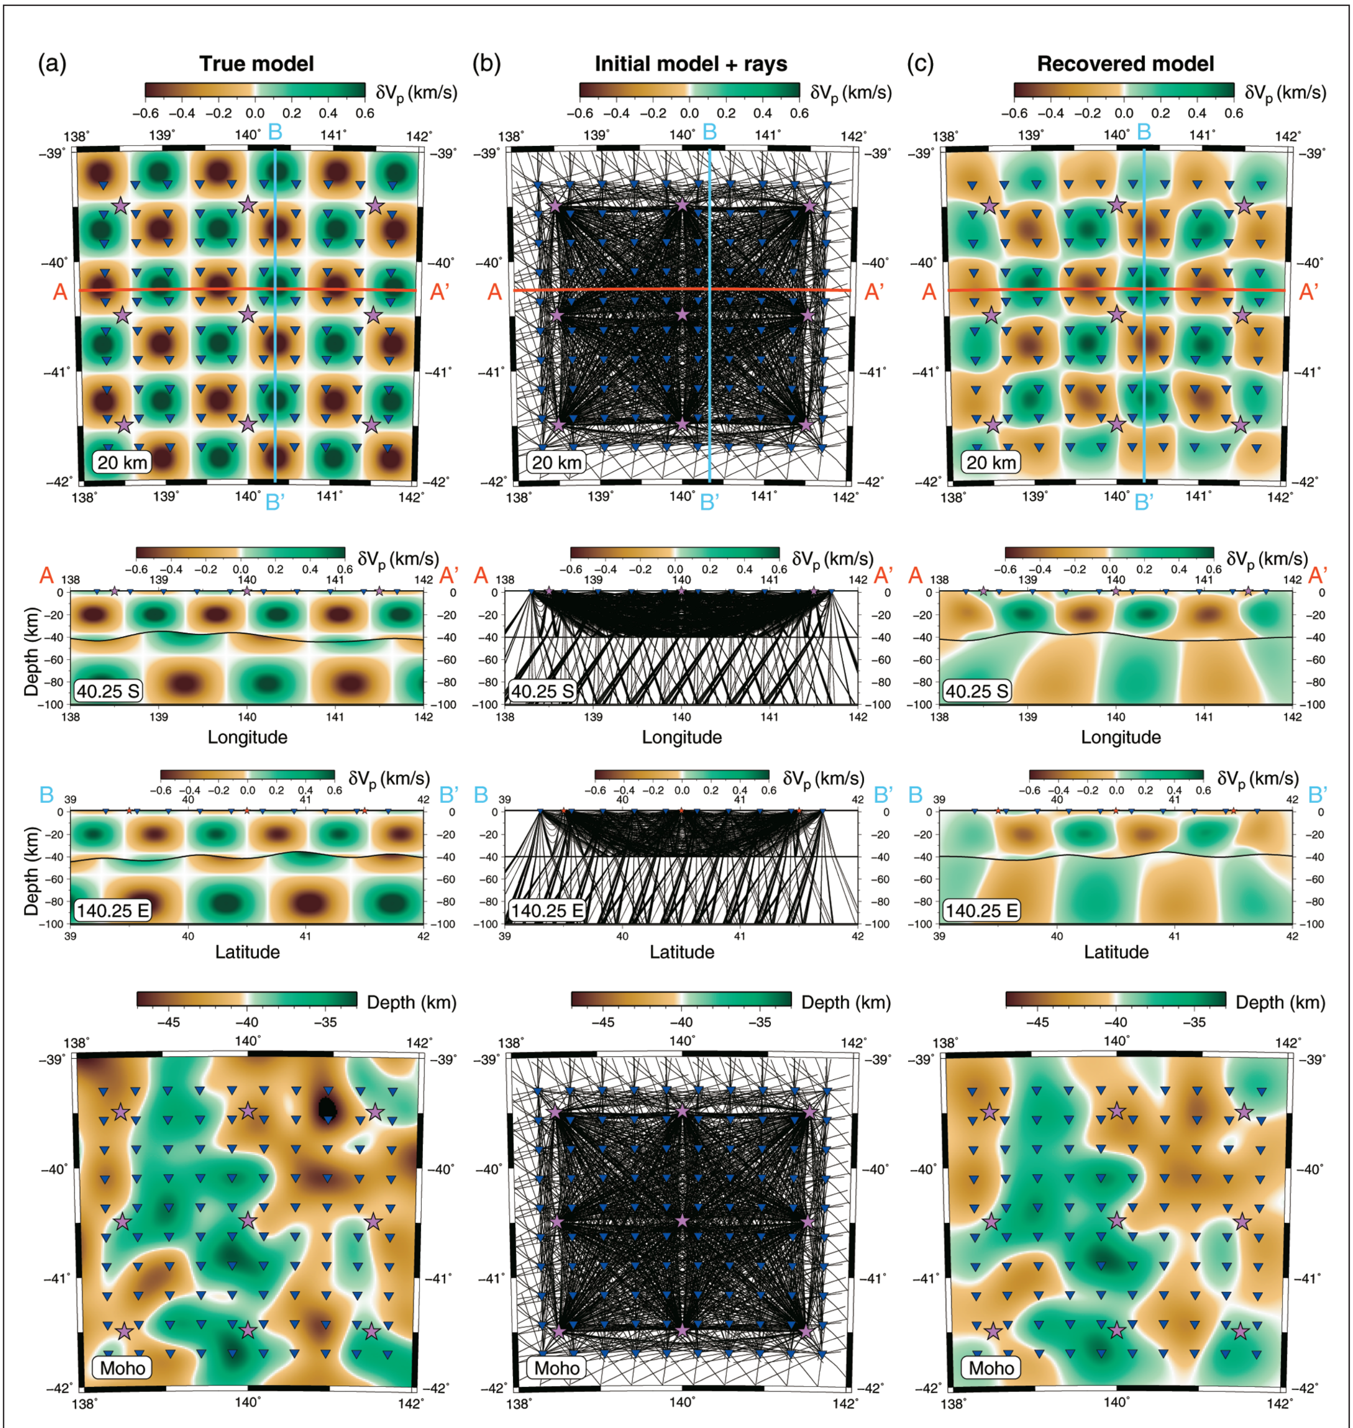

Fig. 7. Imaging results from the combined inversion of synthetic wide-angle and teleseismic data for the 3D velocity and Moho structure of the lithosphere. (a) True, (b) initial, (c) recovered velocity and interface structure. From top to bottom, plots correspond to a depth slice at 20 km, an E-W slice at 40.25°S, a N-S slice at 140.25°E, and Moho depth. Ray path coverage for the initial model is shown in (b). Magenta stars denote sources and blue triangles denote receivers.

satisfactorily resolved. The complete non-linear inversion process, which constrains 14 736 velocity and interface parameters using 3000 traveltimes, takes approximately 25 minutes on a 1.6 GHz Opteron workstation running Linux. The RMS traveltime misfits of the initial and recovered models are 369 ms and 33 ms respectively, which show that the data are well satisfied by the solution. This could possibly be improved slightly by applying more iterations or increasing the density of the FMM propagation grid (to increase traveltime accuracy), but in practical applications there would be little point because of the noise inherent in all observational datasets, and the assumptions imposed by the parameterisation.

The results summarised in Figure 7 show that the multi-stage FMM can be effectively used in sophisticated tomographic inversion routines that combine multiple data classes in order to image complex structures. In addition to teleseismic and wide-angle datasets, coincident reflection and local earthquake data may also be incorporated (hypocentre, velocity and interface parameters could all be inverted for simultaneously), and multi-layered structures that include features such as layer pinch-outs could be imaged provided there is sufficient data coverage. The flexibility offered by a spherical coordinate representation means that the scale of application need not be limited to small regions of the Earth; however, at this stage the scheme is not well suited to global tomography because periodicity is not accounted for, and the degeneracy of the spherical coordinate system at the poles is not recognised. This could be rectified without having to substantially change the code.

CONCLUSIONS

In this paper, a brief description of a new multi-stage FMM scheme for tracking phases in 3D heterogeneous layered media is given, along with several examples that demonstrate its flexibility and robustness in realistic problems. These include: tracking arrivals associated with local earthquake, teleseismic, reflection, and refraction sources through a complex subduction zone; accurately relocating a group of earthquake hypocentres in the presence of significant lateral heterogeneity; and imaging 3D crust and upper mantle structure (including Moho geometry) by combining wide-angle and teleseismic datasets in a single tomographic inversion. Other problems, such as seismic migration, continental-scale body wave tomography, and the imaging of mantle discontinuities using P-S phase conversions, could also be addressed using the new scheme. In many of these areas, ray tracing is still the default tool used for computing traveltimes, but this may soon change with the introduction of grid-based schemes such as the multi-stage FMM, which offer far greater robustness, and in many cases improved efficiency and flexibility.

The source code for the 3D multi-stage FMM scheme used in this paper is freely available for general use from <http://rses.anu.edu.au/seismology/fmmcode>, along with detailed instructions and examples.

ACKNOWLEDGMENTS

This work was funded by Australian Research Council (ARC) Discovery Project DP0451133. N. Rawlinson is supported by ARC Discovery Project DP0556282. All figures were created using the freeware software packages xfig and GMT. D. Brown, M. Purss and an anonymous reviewer are thanked for their constructive comments on an earlier version of the manuscript.

REFERENCES

Aki, K., Christoffersson, A., and Husebye, E.S., 1977, Determination of the three-dimensional seismic structure of the lithosphere: *Journal of Geophysical Research,* **82,** 277–296.

Ankeny, L.A., Braile, L.W., and Olsen, K.H., 1986, Upper crustal structure beneath Jemez Mountains volcanic field, New Mexico, determined by three-dimensional simultaneous inversion of seismic refraction and earthquake data: *Journal of Geophysical Research,* **91,** 6188–6198.

Bai, C.-Y., and Greenhalgh, S., 2005, 3D multi-step travel time tomography: Imaging the local, deep velocity structure of Rabaul volcano, Papua New Guinea: *Physics of the Earth and Planetary Interiors,* **151,** 259–275.

Cao, S., and Greenhalgh, S., 1994, Finite-difference solution of the eikonal equation using an efficient, first-arrival, wavefront tracking scheme: *Geophysics,* **59,** 632–643.

Chopp, D.L., 2001, Some improvements of the fast marching method: *SIAM Journal on Scientific Computing,* **23,** 230–244.

de Kool, M., Rawlinson, N., and Sambridge, M., 2006, A practical grid based method for tracking multiple refraction and reflection phases in 3D heterogeneous media: *Geophysical Journal International,* **167,** 253–270.

Eberhart-Phillips, D., 1986, Three-dimensional velocity structure in northern California coast ranges from inversion of local earthquake arrival times: *Bulletin of the Seismological Society of America,* **76,** 1025–1052.

Graeber, F.M., and Asch, G., 1999, Three-dimensional models of P wave velocity and P-to-S velocity ratio in the southern central Andes by simultaneous inversion of local earthquake data: *Journal of Geophysical Research,* **104,** 20237–20256.

Graeber, F.M., Houseman, G.A., and Greenhalgh, S.A., 2002, Regional teleseismic tomography of the western Lachlan Orogen and the Newer Volcanic Province, southeast Australia: *Geophysical Journal International*, **149**, 249–266.

Hole, J.A., Clowes, R.M., and Ellis, R.M., 1992, Interface inversion using broadside seismic refraction data and three-dimensional traveltime calculations: *Journal of Geophysical Research,* **97,** 3417–3429.

Hole, J.A., and Zelt, B.C., 1995, 3-D finite-difference reflection travel times: *Geophysical Journal International*, **121**, 427–434.

Humphreys, E.D., and Clayton, R.W., 1990, Tomographic image of the Southern California mantle: *Journal of Geophysical Research*, **95**, 19725–19746.

Julian, B.R., and Gubbins, D., 1977, Three-dimensional seismic ray tracing: *Journal of Geophysics*, **43**, 95–113.

Kennett, B.L.N., Sambridge, M.S., and Williamson, P.R., 1988, Subspace methods for large scale inverse problems involving multiple parameter classes: *Geophysical Journal-Oxford*, **94**, 237–247.

Kennett, B.L.N., Engdahl, E.R., and Buland, R., 1995, Constraints on seismic velocities in the Earth from travel times: *Geophysical Journal International*, **122**, 108–124.

Kennett, B.L.N., 2004, Consistency regions in non-linear inversion: *Geophysical Journal International,* **157,** 583–588.

Moser, T.J., 1991. Shortest path calculation of seismic rays: *Geophysics,* **56,** 59–67.

Nakanishi, I., and Yamaguchi, K., 1986. A numerical experiment on nonlinear image reconstruction from first-arrival times for two-dimensional island arc structure: *Journal of Physics of the Earth,* **34,** 195–201.

Parsons, T., and Zoback, M.L., 1997, Three-dimensional upper crustal velocity structure beneath San Francisco Peninsula, California: *Journal of Geophysical Research,* **102,** 5473–5490.

Pereyra, V., Lee, W.H.K., and Keller, H.B., 1980, Solving two-point seismic-ray tracing problems in a heterogeneous medium: *Bulletin of the Seismological Society of America,* **70,** 79–99.

Podvin, P., and Lecomte, I., 1991, Finite difference computation of traveltimes in very contrasted velocity models: a massively parallel approach and its associated tools: *Geophysical Journal International,* **105,** 271–284.

Popovici, A.M., and Sethian, J.A., 2002, 3-D imaging using higher order fast marching traveltimes: *Geophysics*, **67**, 604–609.

Qian, J., and Symes, W.W., 2002, An adaptive finite-difference method for traveltimes and amplitudes: *Geophysics,* **67,** 167–176.

Qin, F., Luo, Y., Olsen, K.B., Cai, W., and Schuster, G.T., 1992, Finite-difference solution of the eikonal equation along expanding wavefronts: *Geophysics,* **57,** 478–487.

Rawlinson, N., Houseman, G.A., and Collins, C.D.N., 2001, Inversion of seismic refraction and wide-angle reflection traveltimes for 3-D layered crustal structure: *Geophysical Journal International*, **145**, 381–401.

Rawlinson, N., and Sambridge, M., 2003, Seismic traveltime tomography of the crust and lithosphere: *Advances in Geophysics*, **46**, 81–198.

Rawlinson, N., and Sambridge, M., 2004a, Wave front evolution in strongly heterogeneous layered media using the fast marching method: *Geophysical Journal International*, **156,** 631–647.

Rawlinson, N., and Sambridge, M., 2004b, Multiple reflection and transmission phases in complex layered media using a miltistage fast marching method: *Geophysics,* **69,** 1338–1350.

Rawlinson, N., and Sambridge, M., 2005, The fast marching method: an effective tool for tomographic imaging and tracking multiple phases in complex layered media: *Exploration Geophysics*, **36,** 341–350.

Rawlinson, N., Reading, A.M., and Kennett, B.L.N., 2006, Lithospheric structure of Tasmania from a novel form of teleseismic tomography: *Journal of Geophysical Research,* **111,** doi:10.1029/2005JB003803.

Riahi, M.A., and Juhlin, C., 1994, 3-D interpretation of reflected arrival times by finite-difference techniques, *Geophysics,* **59,** 844–849.

Sambridge, M.S., and Kennett, B.L.N., 1990, Boundary value ray tracing in a heterogeneous medium: A simple and versatile algorithm: *Geophysical Journal International*, **101**, 157–168.

Sato, T., Kosuga, M., and Tanaka, K., 1996, Tomographic inversion for P wave velocity structure beneath the northeastern Japan arc using local and teleseismic data: *Journal of Geophysical Research,* **101,** 17597–17615.

Sethian, J.A., 1996, A fast marching level set method for monotonically advancing fronts: *Proceedings of the National Academy of Sciences*, **93**, 1591–1595.

Sethian, J.A., 1999, *Level set methods and fast marching methods*: Cambridge University Press.

Sethian, J.A., and Popovici, A.M., 1999, 3-D traveltime computation using the fast marching method: *Geophysics*, **64,** 516–523.

Um, J., and Thurber, C., 1987, A fast algorithm for two-point seismic ray tracing: *Bulletin of the Seismological Society of America*, **77**, 972–986.

Vidale, J.E., 1988, Finite difference calculations of traveltimes: *Bulletin of the Seismological Society of America*, **78**, 2062–2076.

Vidale, J.E., 1990, Finite-difference calculations of traveltimes in three dimensions: *Geophysics,* **55,** 521–526.

Walck, M.C., 1988, Three-dimensional V_P/V_S variations for the Coso region, California: *Journal of Geophysical Research,* **93,** 2047–2052.

Zelt, C.A., and Barton, P.J., 1998, Three-dimensional seismic refraction tomography: A comparison of two methods applied to data from the Faeroe Basin: *Journal of Geophysical Research,* **103,** 7187–7210.