

The fast marching method: an effective tool for tomographic imaging and tracking multiple phases in complex layered media

Nicholas Rawlinson Malcolm Sambridge

Key Words: fast marching method, tomography, ray tracing, teleseismic, seismic wave propagation, refracted wave, reflected wave

ABSTRACT

The accurate prediction of seismic traveltimes is required in many areas of seismology, including the processing of seismic reflection profiles, earthquake location, and seismic tomography at a variety of scales. In this paper, we present two seismic applications of a recently developed grid-based numerical scheme for tracking the evolution of monotonically advancing interfaces, via finite-difference solution of the eikonal equation, known as the fast marching method (FMM). Like most other practical grid-based techniques, FMM is only capable of locating the first-arrival phase in continuous media; however, its combination of unconditional stability and rapid computation make it a truly practical scheme for velocity fields of arbitrary complexity.

The first application of FMM that we present focuses on the prediction of multiple reflection and refraction phases in complex 2D layered media. By treating each layer that the wavefront enters as a separate computational domain, we show that sequential application of FMM can be used to track phases comprising any number of reflection and transmission branches in media of arbitrary complexity. We also show that the use of local grid refinement in the source neighbourhood, where wavefront curvature is high, significantly improves the accuracy of the scheme with little extra computational expense.

The second application of FMM that we consider is in the context of 3D teleseismic tomography, which uses relative traveltimes from distant earthquakes to image wavespeed variations in the Earth's crust and upper mantle beneath a seismic array. Using teleseismic data collected in Tasmania, we show that FMM can rapidly and robustly calculate two-point traveltimes from an impinging teleseismic wavefront to a receiver array located on the surface, despite the presence of significant lateral variations in wavespeed in the intervening crust and upper mantle. Combined with a rapid subspace inversion method, the new FMM based tomographic scheme is shown to be extremely efficient and robust.

INTRODUCTION

One of the classic problems in seismology is to accurately and robustly predict the traveltimes and path of seismic energy between two points within a laterally heterogeneous 2D or 3D medium. Traditionally, this has been solved using geometric ray tracing based on a shooting or bending approach. Shooting methods of ray tracing (e.g., Julian and Gubbins, 1977; Sambridge and Kennett, 1990; Rawlinson et al., 2001) formulate the ray equation as an initial value problem that allows a complete ray to be traced if the source trajectory is specified. The boundary value problem of locating the required two-point path is then solved using an iterative update procedure. The bending method of ray tracing (e.g., Julian and Gubbins, 1977; Um and Thurber, 1987; Grechka and McMechan, 1996) iteratively adjusts the geometry of an initial arbitrary path that joins source and receiver until it becomes a true ray path (i.e., it satisfies Fermat's principle).

The principal drawbacks of ray tracing are related to robustness, speed, and ray selection. In the presence of even small velocity variations, both shooting and bending methods may fail to converge; this lack of robustness increases with the complexity of the medium. Ray tracing can also be a time-consuming process, particularly in the presence of a large number of sources and/or receivers. The final difficulty, that of ray selection, results from the potential existence of multiple two-point paths. Both shooting and bending methods do not necessarily converge to a global-minimum solution (i.e., the first arrival), and often it can be difficult to ascertain which arrival has been located.

A more recently developed and increasingly popular class of method, particularly in the exploration industry, for predicting traveltimes in complex media is to seek finite-difference solutions to the eikonal equation throughout a gridded velocity field (e.g., Vidale, 1988; van Trier and Symes, 1991; Hole and Zelt, 1995; Buske and Kästner, 2004). Although this class of scheme is restricted to locating first arrivals only, the complete traveltimes field can usually be computed extremely rapidly, which allows two-point traveltimes, ray paths, and wavefront geometry to be easily extracted. The main drawback of finite-difference eikonal solvers is that they often suffer from stability problems; in particular, the progressive integration of traveltimes along an expanding square, which is commonly used to compute the traveltimes field, has the potential to breach causality in the presence of large velocity gradients (Qin et al., 1992). However, recent developments with essentially non-oscillatory (ENO) finite difference schemes (e.g., Kim and Cook, 1999) have helped to address this shortcoming.

The problem of calculating traveltimes to every point on a grid can be posed in terms of tracking the evolution of a monotonically advancing interface (e.g., a first-arrival seismic wavefront) that propagates throughout the medium from the source. The need to track an advancing interface is not limited to seismic wavefronts; many other areas of science require this problem to be solved. A recently introduced technique called the fast marching method or FMM (Sethian, 1996; Sethian and Popovici, 1999) was developed

Research School of Earth Sciences
Australian National University
Canberra ACT 0200
Australia
Phone: (02) 6125 0339
Facsimile: (02) 6257 2737
Email: nick@rses.anu.edu.au

Presented at the 17th ASEG Geophysical Conference & Exhibition,
August 2004.
Revised manuscript received 9 June, 2005.

with this in mind. To date, FMM has been applied to a wide variety of problems including optimal path planning, medical imaging, geodesics, and photolithographic development (Sethian, 1999, 2001). In seismology, FMM has been used in the migration of coincident reflection profiles (Popovici and Sethian, 2002).

FMM is a grid-based numerical algorithm which tracks an evolving interface along a narrow band of nodes that are updated by solving the eikonal equation using upwind entropy satisfying finite-difference approximations to the gradient vector. The scheme distinguishes itself from other eikonal solvers by combining unconditional stability (i.e., it always converges to the true solution as grid spacing is decreased) with rapid computation. The unconditional stability of FMM comes from properly addressing the development and propagation of gradient discontinuities in the evolving wavefront.

In this paper, we briefly describe the basic FMM method before demonstrating its use in two practical seismic applications. The first application involves tracking wavefronts composed of any number of reflection and transmission branches in strongly heterogeneous layered media. In order to track this class of multiple, we treat each layer that the wavefront enters as an independent computational domain. Thus, a wavefront is propagated through a layer until it impinges on all points of an interface. From here, a transmitted or refracted branch can be tracked by re-initialising FMM from the narrow band of interface nodes into the adjacent layer, and a reflected branch can be obtained by re-initialising FMM in the incident layer. Wavefronts composed of any number of reflection and transmission branches can therefore be assembled by using this multi-stage approach. Four examples are presented to demonstrate the speed, accuracy, and stability of this new scheme for calculating later-arriving phases.

The second application we present is in 3D tomographic imaging using teleseismic sources. This class of problem typically requires thousands of traveltimes to be predicted in the presence of significant lateral variations in wavespeed; as mentioned previously, conventional ray tracing methods are not always robust in such circumstances. The tomographic imaging method that we develop uses FMM to solve the forward problem of predicting the traveltime residual pattern for each source, and a subspace inversion scheme to solve the inverse problem. In order to account for the non-linearity of the inverse problem, the forward and inverse steps are applied iteratively. To verify the efficiency, robustness, and practicality of the new method, we apply it both to synthetic and observational data.

THE FAST MARCHING METHOD

The basic FMM method for continuous velocity media is briefly described below; for more details, refer to Sethian (1996), Sethian (1999), and Rawlinson and Sambridge (2004). The eikonal equation, which governs the propagation of seismic waves in the high-frequency limit, may be written:

$$|\nabla_x T| = s(\mathbf{x}), \quad (1)$$

where ∇_x is the gradient operator, T is traveltime, and $s(\mathbf{x})$ is slowness as a function of position \mathbf{x} . A significant obstacle for finite-difference methods that seek to solve the eikonal equation for the first-arrival traveltime field is that the wavefront may be discontinuous in gradient. This occurs when the wavefront self-intersects (multi-pathing) but later-arriving information is discarded. The eikonal equation cannot be easily solved in the presence of gradient discontinuities because the equation itself requires $\nabla_x T$ to be defined. FMM deals with this difficulty by enforcing an entropy condition, stated by Sethian and Popovici

(1999) as “once a point burns, it stays burnt”: information can only be lost (or retained in the absence of multi-pathing) as the wavefront evolves because it can only pass through a point once. The unconditional stability of FMM comes from strict enforcement of this condition.

The entropy satisfying upwind scheme we employ is one that has been used by a number of authors including Sethian and Popovici (1999), Chopp (2001), and Popovici and Sethian (2002), and may be written:

$$\left[\begin{aligned} &\max(D_a^{-x}T, -D_b^{+x}T, 0)^2 \\ &+ \max(D_c^{-y}T, -D_d^{+y}T, 0)^2 \\ &+ \max(D_e^{-z}T, -D_f^{+z}T, 0)^2 \end{aligned} \right]^{1/2} = s_{i,j,k}, \quad (2)$$

where (i, j, k) are Cartesian grid increment variables in (x, y, z) , 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. For example, the first two upwind operators for $D^{-x}T_i$ are:

$$\begin{aligned} D_1^{-x}T_i &= \frac{T_i - T_{i-1}}{\delta x} \\ D_2^{-x}T_i &= \frac{3T_i - 4T_{i-1} + T_{i-2}}{2\delta x} \end{aligned} \quad (3)$$

where δx is the grid spacing in x . Which operator is used in equation (2) depends on the availability of upwind traveltimes and the maximum order allowed. First-order schemes only use D_1 operators, second-order schemes preferentially use D_2 operators but revert to D_1 if T_{i-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. Successful implementation of this scheme requires that the order in which nodes are updated be consistent with the direction of flow of information; that is, from smaller to larger values of T . FMM achieves this by systematically constructing traveltimes in a downwind fashion from known values upwind using a narrow-band approach. The narrow band concept is illustrated in Figure 1. All grid points are labelled as either *Alive*, *Close* or *Far*. *Alive* points lie upwind of the narrow band and have correct traveltime values; *Close* points lie within the narrow band and have trial values calculated using equation 2 with *Alive* points only; *Far* points




Fig. 1. Principle of the narrow band method for tracking the first arrival wavefront. See text for more details.




Fig. 2. Implementing grid refinement in the source neighbourhood. When the narrow band (thick grey line) reaches the boundary of the refined grid, it is mapped onto the coarse grid (denoted by triangles) before continuing to evolve.




Fig. 3. Interface nodes (black triangles) are locally sutured to adjacent velocity nodes (grey dots) via a triangular mesh (dashed black line). The thick grey line represents the actual interface.

lie downwind of the narrow band and have no traveltime values calculated. The narrow band is evolved by identifying the *Close* point with minimum traveltime (using a heap sort algorithm), tagging it as *Alive*, and then tagging all neighbouring *Far* points as *Close*. Finally, all *Close* points adjacent to the new *Alive* point have their traveltimes updated using equation 2. 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 computational cost scales with grid size in this way is responsible for the overall efficiency of FMM.

Point Source Error Minimisation

A point source is an upwind singularity of the traveltime field and can be a major contributor to FMM traveltime error because of high wavefront curvature and limited (first-order) accuracy in its vicinity. A variety of methods have been proposed to address this problem, which afflicts most grid-based eikonal solvers, including using a spherical grid centred on each source point (Alkhalifah and Fomel, 2001); local grid refinement in the source neighbourhood (Kim and Cook, 1999); adaptive gridding with grid refinement and coarsening based on a posteriori error estimation (Qian and Symes, 2002).

The approach we use for minimising error in the source neighbourhood is based on specifying a finely spaced grid in the vicinity of the source and a coarser grid away from the source, making it similar to the method advocated by Kim and Cook (1999), who use an ENO scheme to solve the eikonal equation. In their approach, the traveltime field is computed along an expanding box; in our case, the traveltime field is computed along an expanding wavefront, which means that the computational front may pass from the coarse grid back into the finer grid if the local grid about the source is defined by a rectangular region. This can produce numerical instabilities (see Sethian, 1999). To overcome this problem, we define a rectangular grid of refined nodes about the source, but downsample the computed traveltime field to the coarse grid spacing when the first node on the edge of the refined grid becomes *Alive* (see Figure 2). As a result, the true edge of the refined grid will conform to the shape of the narrow band, so that information only flows out of the refined grid and never back into it. This approach to local grid refinement ensures that the stability of FMM is not compromised.

APPLICATION TO COMPLEX LAYERED MEDIA

The Multi-Stage FMM

In a layered medium with undulating interfaces, a purely regular grid is no longer suitable for describing structural variations. Instead, we define velocity on a regular grid, and then use an adaptive triangular mesh to suture interface nodes to neighbouring velocity nodes, as illustrated in Figure 3. Interface nodes are defined by the intersection points of the interface (defined by cubic B-splines in our examples) with the cell boundaries of the rectangular velocity grid. Within this framework, the irregular mesh is constructed so that (1) triangles do not span more than one rectangular velocity cell or pass through the interface, and (2) the presence of obtuse triangles is minimised. Equation 2 is used to update grid points in the regular mesh, and a first-order accurate scheme for triangular elements is used to evolve the wavefront through the irregular mesh. Further details of the irregular mesh scheme can be found in Rawlinson and Sambridge (2004).

To understand how FMM can be used to track more than one arrival, consider Figure 4a, which shows a wavefront emanating from a point source and impinging upon an interface. Rather than continuing to propagate through to the adjacent layer, the wavefront is only tracked as far as the interface, which is treated as one of four boundaries of the computational domain. Once all points contained in the layer are *Alive*, the narrow band ceases to exist and the FMM process is terminated. The next step is to track a transmitted or reflected wavefront using the traveltime information

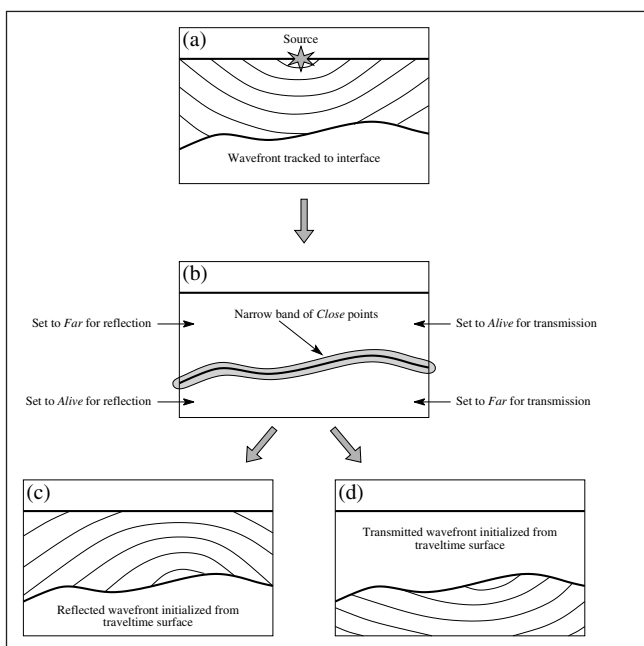


Fig. 4. Principle behind the tracking of reflection and transmission traveltimes. The incident wavefront is tracked to all points on the wavefront, before FMM is reinitialised in the incident (for reflection) or adjacent (for transmission) layer.

we have obtained as a starting point. This can be done by using interface node traveltimes only, as they store sufficient information for first-order accurate estimates of traveltimes to neighbouring nodes. The complete set of interface nodes is therefore used as the starting narrow band for the next FMM stage, as illustrated in Figure 4b. The fact that this narrow band will not in general conform to the shape of the wavefront will not compromise the integrity of the scheme because the wavefront is not permitted to reflect from or transmit through an interface more than once in a single FMM step.

From the initial narrow band described by the set of interface nodes, a reflected wavefront can be tracked by setting all nodes within the incident layer to *Far*, and re-starting FMM (see Figure 4c). A transmitted wavefront is tracked in the same manner, except that only nodes in the adjacent layer are set to *Far*, so that the wavefront continues to propagate into the next layer (see Figure 4d) rather than reflecting back into the incident layer. A transmitted wavefront can also be obtained by passing the wavefront through an interface in a single FMM iteration; in this case, the wavefront may transmit through the interface more than once. The basic approach of partitioning each layer into separate computational domains as the wavefront evolves can be repeated any number of times to track any required phase. Significantly, memory resources are proportional to the number of nodes within a layer, not the total number of nodes traversed by the wavefront. Consequently, the number of bounces and transmissions that can

be tracked is limited by CPU time only.

In practice, the total energy carried by the wavefront decreases monotonically with time, so the modelling of real observations is unlikely to require that phases composed of large numbers of reflection and transmission branches be tracked.

Relatively few previous studies have attempted to compute later-arriving phases using grid-based eikonal solvers. Podvin and Lecomte (1991) show that first-arriving head waves, which often have small amplitude, can be suppressed by making the underlying high-velocity layer a masking layer with zero velocity. This allows the later-arriving phase from an overriding lower-velocity layer to be tracked. Podvin and Lecomte (1991) also find reflections by tracking the first-arrival traveltimes fields from both the source and receiver to the entire interface. Valid reflection points along the interface can then be identified using Fermat's principle of stationary time.

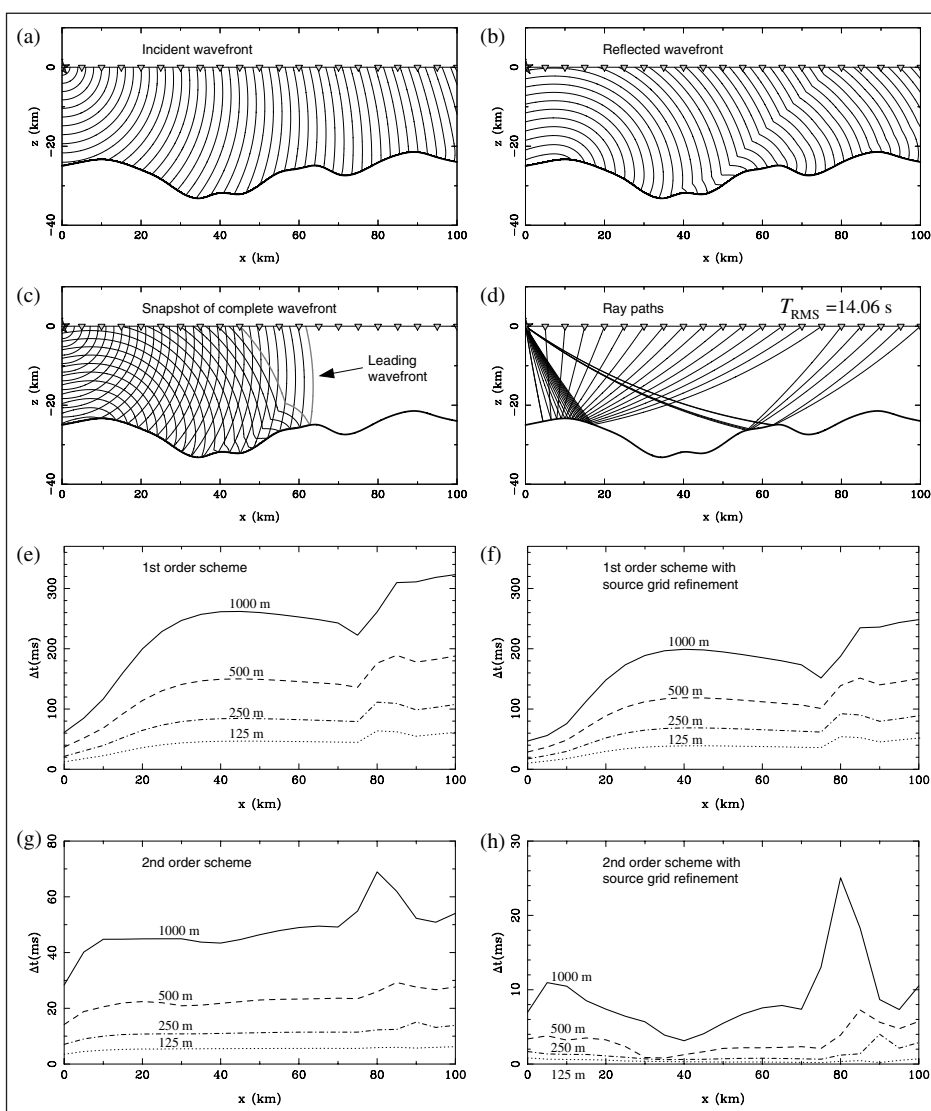


Fig. 5. Tracking of a single reflection arrival through a medium containing a single interface. Wavespeed varies linearly with depth within a layer, so each ray path segment is described by a circular arc. Wavefronts are plotted at 0.4 s intervals in all cases. Receivers are denoted by grey triangles and the source by a dark grey star. (a)–(d) Wavefronts and rays; (e)–(h) traveltime accuracy of four different schemes using four different grid sizes (1000 m, 500 m, 250 m, 125 m). See also Table 1.

Although multiple reflection points can be found for each source-receiver pair, the clear drawback is that a traveltimes field must be computed for each source and each receiver. Hole and Zelt (1995) overcome this problem by using a local plane wave approximation and Snell's Law in the neighbourhood of an interface to track the reflecting wavefront. However, strong wavefront or reflector curvature will degrade the accuracy of this scheme.

The tracking of later arrivals using eikonal methods has also been considered in the context of seismic migration. Zhao et al. (1998) generalise the approach of Podvin and Lecomte (1991), to improve Kirchhoff diffraction mapping using a limited class of later arrivals. Manuel and Uren (2000) perform pre-stack depth migration of seismic multiples by extending the work of Zhao et al. (1998) to permit the calculation of reflections. This is achieved by implementing the so-called generalised exploding reflector model or GERM (see also Manuel et al., 2001). The principle of GERM is to explode a series of equally spaced point sources along a reflector with explosion times equal to the incident time of the impinging wavefront. The correct shape of the reflected wavefront is then obtained by selecting only the first-arrival from all exploding sources using Huygens' principle. Although the implementation of GERM is quite different to our multi-stage FMM scheme, the principle behind the evolution of a reflecting wavefront from an interface is similar.

Examples

Four examples are presented to assess the viability of FMM for predicting the traveltimes of phases comprising multiple reflection and transmission branches in complex 2D layered media. In these examples, the velocity fields and interfaces are described using cubic B-spline functions, which are controlled by separate sets of velocity and interface vertices. Discrete sampling of these functions, which are continuous in curvature, is easily accomplished for any specified resolution. When grid refinement in the source neighbourhood is applied, we use a refined grid that extends 10 coarse node points horizontally and vertically from the source and has a fivefold decrease in node spacing. We found that increasing the area of the refined grid and decreasing the node spacing has only a marginal effect on the accuracy of the computed traveltimes.

The first example (see Figure 5) consists of a single reflection from an undulating surface that lies within a velocity field with a constant vertical gradient of 0.04 s^{-1} . This phase is tracked by invoking FMM twice: once for the incident wavefront (Figure 5a), and once for the reflected wavefront (Figure 5b). A snapshot of the complete wavefront, which is obtained by stitching together

matching isochrons from both the incident and reflected traveltimes fields, is shown in Figure 5c. Source receiver ray paths (Figure 5d) are obtained a posteriori by following the direction of steepest descent from each receiver through the traveltimes field to the source.

Because of the undulating interface, the velocity model of Figure 5 is too complex to allow analytic solutions. Therefore, in order to analyse the accuracy of FMM, highly accurate two-point traveltimes are obtained using a shooting method of ray tracing (Rawlinson et al., 2001). This scheme uses analytic solutions in the presence of constant velocity gradients, and allows the accuracy of ray-interface intersection points to be controlled, in addition to the distance between receiver and ray end-point. In the following comparison, the ray interface intersection points are accurate to 0.5 mm and the ray-receiver intersection points are accurate to 5 mm. Such small values will result in ray tracing traveltimes that for our purposes can be considered exact.

Figures 5e–5h show error plots for four different FMM schemes with velocity grid spacings of 1000 m, 500 m, 250 m and 125 m. The error estimates are equal to the difference between the FMM solution and the ray tracing solution at each of the 21 receivers. Table 1 shows the RMS errors for the four different schemes with the four grid sizes that are plotted, but also includes grid sizes of 62.5 m and 31.25 m (not included in Figures 5e–5h for clarity). CPU times (on a 1.6 GHz Opteron PC, running Linux) for each of these calculations are also shown. The first-order scheme (Figure 5e) exhibits first-order convergence; the only slight increase in accuracy achieved by including grid refinement about the source (Figure 5f) suggests that near-source errors are not principally responsible for overall error. The second-order scheme (Figure 5g) is significantly more accurate than both of the first-order schemes. With the addition of local grid refinement about the source (Figure 5h), the scheme becomes much more accurate again and the convergence improves. This suggests that errors in the source neighbourhood dominate the overall error of the second-order scheme. A comparison of errors and CPU times (Table 1) clearly suggest that the most efficient scheme is the second order scheme with grid refinement about the source. In fact, to achieve the same accuracy with a first-order scheme, approximately three orders of magnitude more computation time is required.

The second example (Figure 6) shows a multiple tracked through a structure composed of two layers overlying a half-space; unlike the previous example, both the interfaces and the layer velocities vary laterally. In order to track this multiple, FMM is invoked four times, as suggested by the schematic of Figure 6a. Each component of the wavefield is shown in Figures 6b–6e. A

Grid spacing (m)	Number of grid points	RMS error (ms)				CPU time (s)			
		1st O.	1st O.R.	2nd O.	2nd O.R.	1st O.	1st O.R.	2nd O.	2nd O.R.
1000	4141	243.1	181.3	48.7	10.1	0.01	0.02	0.01	0.02
500	16281	142.2	111.4	23.2	3.5	0.04	0.05	0.03	0.05
250	64561	81.3	65.9	11.4	1.5	0.17	0.17	0.18	0.19
125	257121	45.7	38.0	5.5	0.5	0.70	0.71	0.76	0.79
62.5	1026241	25.3	21.5	2.7	0.2	3.15	3.23	3.41	3.59
31.25	4100481	13.9	12.0	1.3	0.1	13.67	14.22	15.09	15.58

Table 1. Summary of results involving a single reflection from an undulating interface and no lateral velocity variations within each layer (see Figure 5). 1st O. = first order scheme; 1st O.R. = first order scheme with local grid refinement in the source neighbourhood; 2nd O. = second order scheme; 2nd O.R. = second order scheme with local grid refinement in the source neighbourhood. CPU times are for a 1.6 GHz Opteron PC, running Linux.

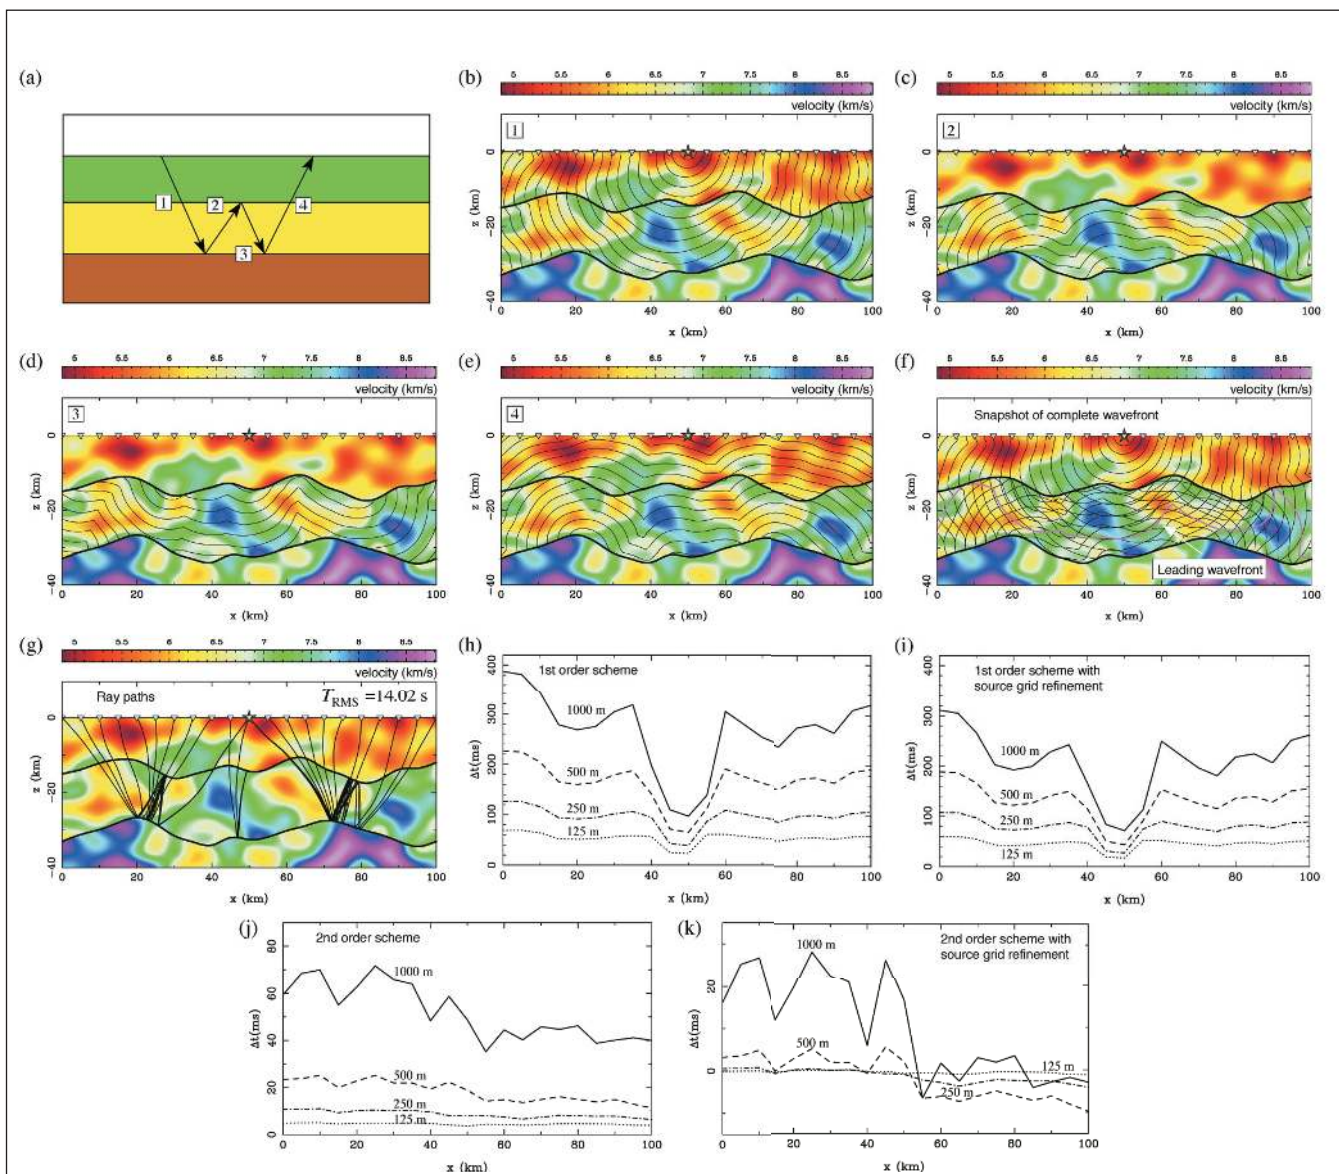


Fig. 6. Tracking a multiple through a layered medium with laterally varying interface structure and velocity. Wavefronts are plotted at 0.4 s intervals in all cases. (a) Structure of the multiple that is tracked; (b)–(e) each of the four branches of the wavefield; (f), (g) snapshot of complete wavefront and corresponding ray paths; (h)–(k) traveltime accuracy of four different schemes using four different grid sizes (1000 m, 500 m, 250 m, 125 m). See also Table 2.

snapshot of the complete wavefield is shown in Figure 6f, and ray paths to all 21 receivers are shown in Figure 6g. It is interesting to note that discontinuities in wavefront curvature pose no problem to the multi-stage FMM. For example, the wavefield in Figure 6d exhibits an evolving discontinuity at about $x = 65$ km, with no sign of instability. In Figure 6e, the same discontinuity has been reflected, and again, there are no apparent signs of instability. Earlier schemes for solving the eikonal equation would have difficulty with this complexity of wavefront evolution.

The accuracy of the multi-stage scheme for the Figure 6 example cannot be analysed in the same way as Figure 5, because the ray-tracing scheme cannot track multiples. In addition, the presence of laterally varying velocity layers means that any ray-tracing scheme would have difficulty achieving high accuracy and guaranteeing the location of the first arrival of the multiple that is tracked. Instead, we appeal to the numerical stability of FMM, as exhibited in the first example, and use the solution calculated with the second-order scheme with grid refinement about the source on a 31.25 m grid as a proxy for the exact solution. As Figures 6h–6k and Table 2 show, the four different schemes exhibit similar

behaviour to the Figure 5 example as the grid size is reduced, and it is clear that the second order scheme with grid refinement in the source neighbourhood offers the best compromise between computational speed and accuracy.

One of the classical problems in seismic refraction and wide-angle reflection analysis is the prediction of observed refraction and reflection branches in the presence of undulating interfaces separating layers containing lateral variations in wavespeed. Although shooting and bending methods of ray tracing are often used to solve this class of problem (e.g., Zelt and Smith, 1992), they can encounter difficulties even when lateral variations are small. For example, refracted rays near the critical angle can be difficult to compute because the variation in ray take-off angle with the location of the ray end-point is highly non-linear. Additional difficulties include the accurate tracking of head waves, and consistent identification of the same branch of a triplicated reflection caused by variations in interface topography. The multi-stage FMM scheme overcomes all of these problems because it computes the complete traveltime field, always tracks the first arrival within each computational domain, and is unconditionally stable.

Figure 7 shows an example using a similar velocity model to Figure 6, except that the lateral variations in structure are slightly more muted and the velocity contrasts across the interfaces are slightly larger. This has been done simply to make the plotting of different traveltimes easier to understand. For the given model, the phases that are typically identified from a refraction section are P (direct arrival), P_1 (refraction from below first interface), P_2 (refraction from below second interface), P_1P (reflection from top interface) and P_2P (reflection from bottom interface), as shown in Figure 7a. The multi-stage FMM scheme can be used to track each of these phases, and construct their complete traveltimes curves, as shown in Figure 7b. Note the

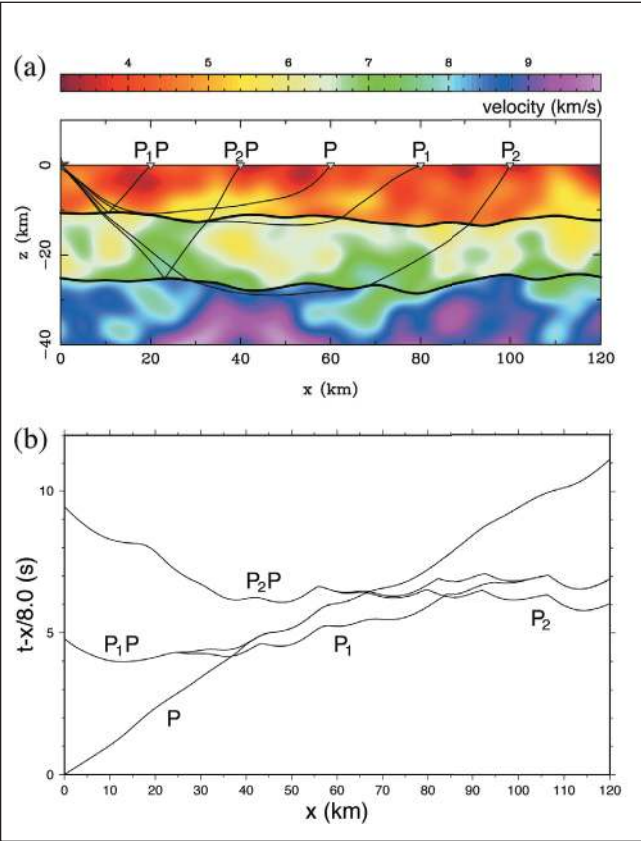


Fig. 7. Predicting the traveltime curves of five different crustal phases. (a) An example ray of each of the five phases; (b) traveltime curves for the five phases shown on a reduced traveltime plot.

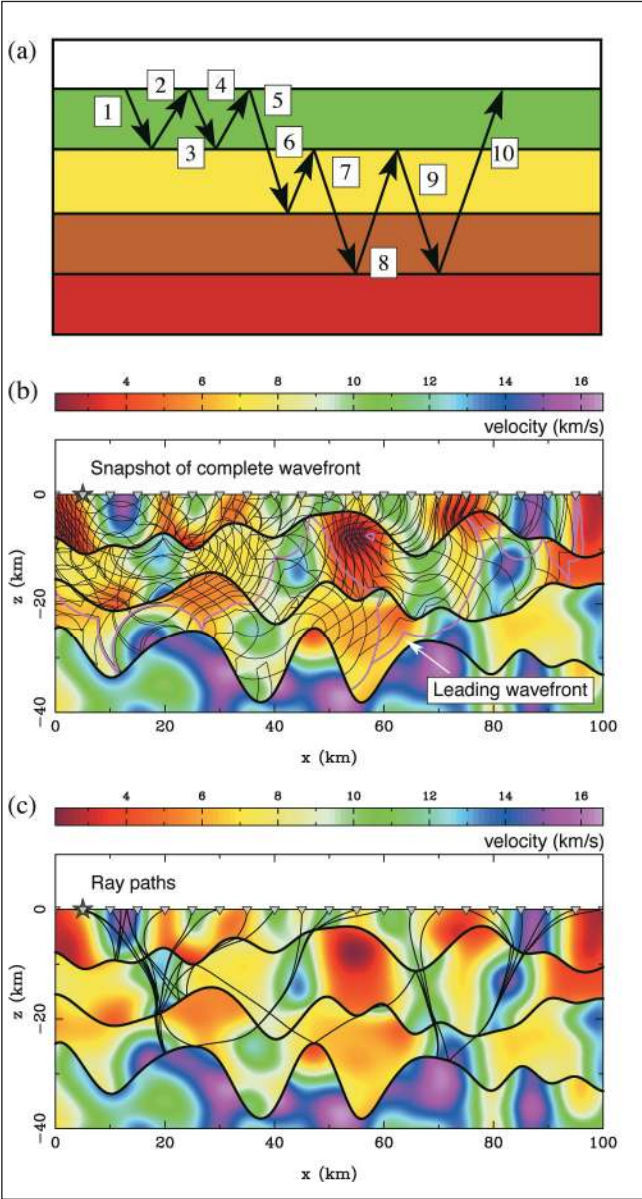


Fig. 8. Tracking a 10-fold multiple through a complex layered medium. (a) Structure of the multiple; (b) snapshot of the complete wavefront (plotted at 0.4 s intervals); (c) ray paths from the source to 21 receivers located on the surface.

Grid spacing (m)	Number of grid points	RMS error (ms)				CPU time (s)			
		1st O.	1st O.R.	2nd O.	2nd O.R.	1st O.	1st O.R.	2nd O.	2nd O.R.
1000	4141	278.4	218.1	53.12	15.4	0.02	0.02	0.03	0.04
500	16281	170.3	135.8	18.9	5.4	0.07	0.08	0.08	0.10
250	64561	98.5	81.0	8.8	2.2	0.28	0.29	0.31	0.33
125	257121	55.7	47.0	4.5	0.6	1.14	1.15	1.26	1.31
62.5	1026241	30.7	26.4	2.2	0.3	5.11	5.30	5.73	5.99
31.25	4100481	16.8	14.6	1.2	***	23.48	24.23	26.36	26.90

*** = Reference traveltime used for error estimate

Table 2. Summary of results involving a four-fold multiple in a medium containing two interfaces and lateral velocity variations within layers (see Figure 6). 1st O. = first order scheme; 1st O.R. = first order scheme with local grid refinement in the source neighbourhood; 2nd O. = second order scheme; 2nd O.R. = second order scheme with local grid refinement in the source neighbourhood. CPU times are for a 1.6 GHz Opteron PC, running Linux.

complexity of each of the two triplications formed by P , P_1 and P_1P , and P_1 , P_2 and P_2P ; ray tracing schemes would be hard pressed to accurately predict the traveltime curves in these regions. Additional phases, such as multiples, can easily be added to the traveltime curve plot, and increasing the complexity of the medium poses no difficulty to the multi-stage FMM scheme.

The final example (Figure 8) demonstrates that FMM is capable of tracking phases composed of any number of refraction and reflection branches in media of virtually any complexity. Figure 8a schematically illustrates the 10 branches of a complex phase that we attempt to track. Figures 8b and 8c show the almost pathological variations in interface structure (up to 15 km changes in relief over horizontal distances of 10 km) and wavespeed (lateral velocity contrasts as great as 8:1) that the multi-stage FMM has to deal with. Figure 8b shows a snapshot of the evolving wavefront, which exhibits many propagating gradient discontinuities. Source-receiver ray paths are shown in Figure 8c; it would be safe to say that conventional ray tracing schemes would be unable to locate such complex two-point paths. The trade-off between accuracy and CPU time for this example using different schemes and grid sizes is not shown, but exhibits similar behaviour to the two previous examples. An important point to note in Figure 8c is that the ray paths tend to favour the fast region of the model; this is because only the first arrival is tracked within a single computational domain. This means that the energy carried by the wavefront decreases monotonically with time, and that in practice, phases such as the one tracked in Figure 8 would probably not be visible on any seismic record.

The above examples illustrate that FMM can be successfully applied in the presence of complex layered media to track phases composed of any number of reflection and transmission branches. The second order scheme with source grid refinement clearly offers the optimum trade-off between CPU time and accuracy. The computational efficiency and robustness of the multi-stage FMM makes it a valuable tool for the analysis of refraction or coincident reflection data; for instance, to remove multiples or in the forward step of a traveltime inversion for crustal structure.

APPLICATION TO 3D SEISMIC TRAVELTIME TOMOGRAPHY

FMM in Spherical Coordinates

The following example images a region of Tasmania at a scale that requires the Earth's sphericity to be accounted for. Therefore, the scheme for updating grid points defined by equations (2) and (3) needs to be modified for spherical coordinates. In practice, this can be done most simply by reformulating the operators defined in equation (3) and keeping equation (2) unchanged. In spherical coordinates, the first-order upwind difference operators for D_i^{-r} , $D_j^{-\theta}$, $D_k^{-\phi}$ are:




Fig. 9. Illustration of the teleseismic tomography principle. A local model volume containing 3D variations in velocity is defined beneath the receiver array. Outside this region, the Earth is assumed to be spherically symmetric.

$$\begin{aligned} D_1^{-r}T_i &= \frac{T_i - T_{i-1}}{\delta r} \\ D_1^{-\theta}T_j &= \frac{T_j - T_{j-1}}{r\delta\theta} \\ D_1^{-\phi}T_k &= \frac{T_k - T_{k-1}}{r\sin\theta\delta\phi} \end{aligned} \quad (4)$$

where δr , $\delta\theta$, $\delta\phi$ are the grid spacing in depth, co-latitude and longitude, respectively. The equivalent expressions for the second order upwind operators are:

$$\begin{aligned} D_2^{-r}T_i &= \frac{3T_i - 4T_{i-1} + T_{i-2}}{2\delta r} \\ D_2^{-\theta}T_j &= \frac{3T_j - 4T_{j-1} + T_{j-2}}{2r\delta\theta} \\ D_2^{-\phi}T_k &= \frac{3T_k - 4T_{k-1} + T_{k-2}}{2r\sin\theta\delta\phi} \end{aligned} \quad (5)$$

Use of these operators makes converting from Cartesian to spherical coordinates a relatively straightforward task, as the remainder of the scheme remains essentially unchanged.

Example

The following example makes use of TIGGER data (Rawlinson and Kennett, 2004) collected by the Seismology Group at the Research School of Earth Sciences, Australian National University. The TIGGER experiment is a multifaceted seismic experiment, the field component of which was carried out in Tasmania and southern Victoria in 2001/2002. As part of this experiment, an array of 64 short period and 8 broadband seismometers with a 15 km spacing was deployed in northern Tasmania for a period of 5 months to record distant (teleseismic) earthquakes. To date, relative traveltime residuals (Rawlinson and Kennett, 2004) from 100 events from regions such as Indonesia, Japan, the Philippines, Papua New Guinea, New Zealand, Fiji, South Sandwich Islands etc. have been picked. The complete data set comprises nearly 6000 P-wave traveltime residuals.

3D teleseismic traveltime tomography involves mapping the traveltime residuals as lateral variations in seismic wavespeed in the Earth's crust and upper mantle. Usually, this is done by specifying a model region beneath the seismic array within which 3D variations in wavespeed are permitted (e.g., Humphreys and Clayton, 1990; Steck et al., 1998; Graeber et al., 2002); outside this region, the Earth is assumed to be spherically symmetric. This allows traveltimes for any of the major global phases to be rapidly computed to the edge of the model region using a reference model such as *ak135* (Kennett et al., 1995). However, traveltimes within the local model beneath the array are usually computed using 3D ray tracing (e.g., vanDecar et al., 1995). This can be very time consuming for large datasets involving many source-receiver pairs, and if iterative non-linear inversion methods are used to adjust the model parameters to satisfy the data, then the forward problem of predicting traveltime residual patterns must be repeated many times. In addition, methods such as ray tracing are not robust, and in the presence of even mild lateral heterogeneities, it is possible that a small but significant subset of traveltimes will not be predicted.

Here, for the first time, we examine the use of FMM in 3D seismic traveltime tomography. By combining computational stability and robustness, FMM guarantees the rapid prediction of all observed traveltimes. Unlike the previous examples which examined refraction and reflection phases in crustal models, FMM is not initiated from a point source, because distant

traveltimes are computed to the bottom face of the local model using a global reference model (*ak135* in this case), as shown schematically in Figure 9. Therefore, the same principle that we used previously to initiate FMM from an interface can be applied, and no local grid refinement is required. We use cubic B-spline functions to parameterise the local model volume; discretisation of these continuous functions to a 3D mesh of points through which the traveltime field is calculated is straightforward and computationally efficient.

The inverse problem in tomography is to adjust the values of the model parameters in order to satisfy the data, subject to regularisation constraints. In our case, the data are relative traveltime residual patterns, and the unknowns are the grid of vertices that control the pattern of the B-spline velocity field. In order to solve this problem, we formulate the inverse problem as an optimisation problem, where an objective function $S(\mathbf{m})$, where \mathbf{m} is the vector of unknowns, is minimised. The objective function is defined in this case by:

$$S(\mathbf{m}) = (\mathbf{g}(\mathbf{m}) - \mathbf{d}_{\text{obs}}) \mathbf{C}_d^{-1} (\mathbf{g}(\mathbf{m}) - \mathbf{d}_{\text{obs}})^T + \varepsilon (\mathbf{m} - \mathbf{m}_0) \mathbf{C}_m^{-1} (\mathbf{m} - \mathbf{m}_0)^T, \quad (6)$$

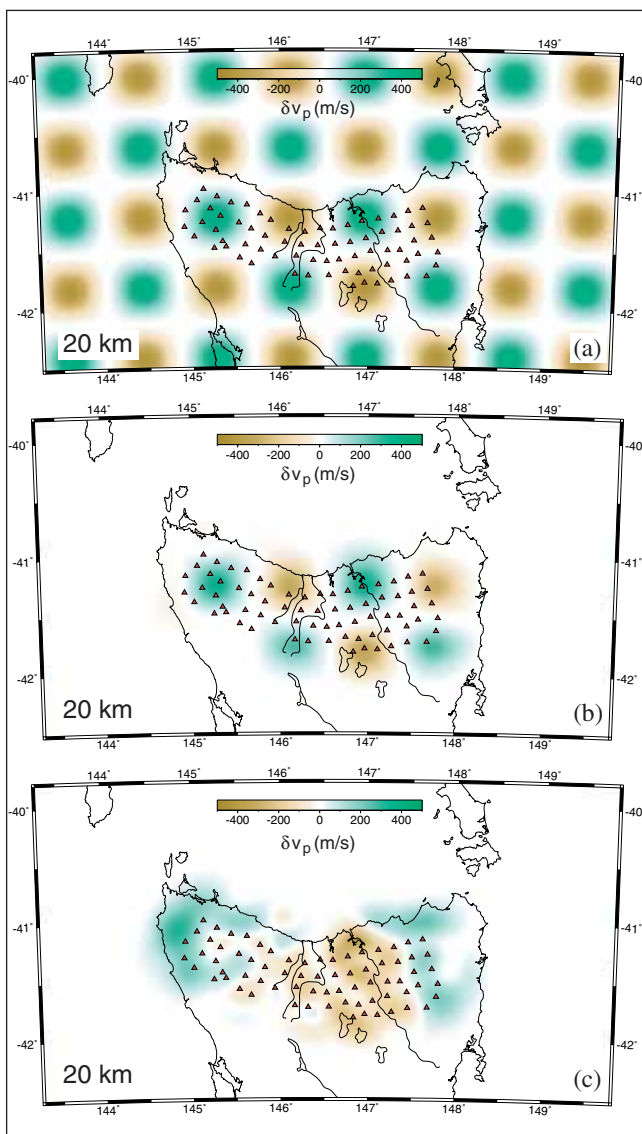


Fig. 10. Horizontal slices through (a) the synthetic checkerboard model; (b) recovered checkerboard model; (c) the TIGGER solution model, at 20 km depth. Velocity variations are plotted as perturbations from *ak135*.

where \mathbf{d}_{obs} are the observed traveltime residuals, $\mathbf{g}(\mathbf{m})$ are the predicted residuals, \mathbf{m}_0 is the reference model (*ak135*), \mathbf{C}_d^{-1} is the a priori data covariance matrix, \mathbf{C}_m^{-1} is the a priori model covariance matrix, and ε is the damping parameter. The first term on the right-hand side of equation (6) aims to find a model which satisfies the data; the second term is a regularisation term which penalises models which differ too greatly from the initial or reference model (see Rawlinson and Sambridge, 2003, for more details). We minimise equation (6) using an iterative non-linear approach, which successively uses a second order FMM scheme to solve the forward problem of predicting traveltime residuals, and a subspace inversion method (Kennett et al., 1988) to solve the inverse problem of adjusting the model parameters to satisfy the data and regularisation constraints.

The 3D tomographic imaging scheme outlined above is used to construct a preliminary image of the lithosphere beneath northern Tasmania from 5938 traveltime residuals. The 3D model is parameterised using 9996 velocity nodes (approximately 15 km separation in each dimension), and six iterations of a 10-D subspace inversion method are used to minimise the objective function. The complete calculation takes only 8.5 minutes on a 1.6 GHz Opteron workstation, running Linux. This is extremely rapid for the iterative non-linear solution of a 3D tomographic imaging problem, and makes practical the detailed analysis of solution robustness based on repeated tests.

Figures 10 and 11 show two slices through the solution model. In addition to the model derived from the TIGGER observations,

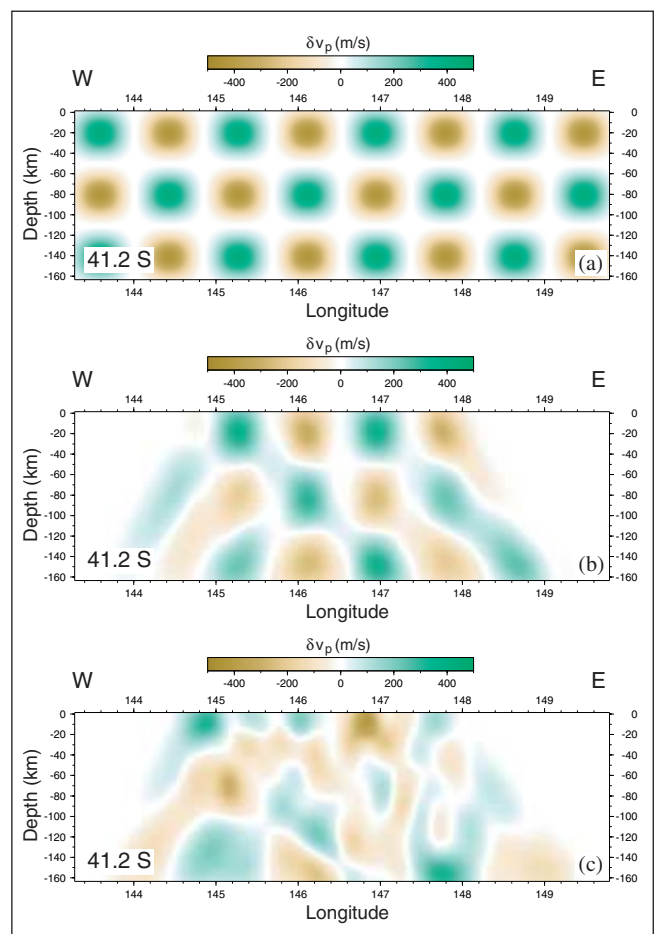


Fig. 11. Vertical E-W slices through (a) the synthetic checkerboard model; (b) recovered checkerboard model; (c) the TIGGER solution model, at 41.2° S. Velocity variations are plotted as perturbations from *ak135*.

the results of a synthetic checkerboard test are also shown, which reflect the resolving power of the data. The test works by constructing a synthetic model, using FMM to predict the residual pattern for the distribution of events that produce the observed data, and then applying the tomographic imaging routine to the synthetic data in order to reconstruct the synthetic model. Figure 10a shows a slice through the synthetic checkerboard at 20 km depth, and Figure 10b shows the reconstruction. At this depth, ray coverage outside the horizontal bounds of the array is not significant, which explains why the checkerboard is not recovered in this region. However, the recovery is excellent beneath the array, which both demonstrates that ray coverage is good and that the inversion scheme is working correctly. The solution model obtained from real data is shown in Figure 10c. A vertical E-W slice taken at 41.2° S through the synthetic checkerboard, recovered checkerboard, and TIGGER model are shown in Figures 11a, 11b, and 11c respectively. Although there is some smearing present in Figure 11b, the pattern of synthetic anomalies is generally well recovered.

Interpretation of the TIGGER model will be the subject of a future paper, as we only wish to demonstrate the viability of FMM in the forward step of a 3D tomographic inversion problem. However, the velocity distribution that is recovered shows promise in helping to unravel the deep geology and tectonic evolution of the Tasmanian region. Of particular significance is the distinct E-W pattern of fast – slow – fast anomalies in the deep crust (Figure 10c). Although we have demonstrated the viability of FMM in teleseismic tomography, our results are equally applicable to other classes of tomography, such as local earthquake tomography and refraction tomography.

CONCLUSIONS

FMM is a fast and unconditionally stable grid-based scheme for directly solving the eikonal equation using finite differences. Originally developed in the field of computational mathematics, it has only seen limited application in the seismology community to date. In this paper, we show it to be a highly effective method for determining the traveltimes of phases composed of any number of reflection and transmission branches in complex laterally heterogeneous velocity media. The use of grid refinement in the neighbourhood of a source point and higher order operators results in dramatic improvements in accuracy compared to the original first order scheme, and helps make the multi-stage FMM for layered media an extremely effective tool for seismic refraction and reflection analysis. We also show that the speed and robustness of FMM can be used to great effect in large 3D tomographic inversion problems. In particular, an iterative non-linear teleseismic tomography problem involving nearly 10 000 unknowns and 5938 ray paths is solved in less than 10 minutes on a standard workstation.

ACKNOWLEDGMENTS

This work was supported by Australian Research Council Discovery Project DP0451133. All figures were created using the freeware software packages PGPlot, xfig, and GMT.

REFERENCES

- Alkhalifah, T., and Fomel, S., 2001, Implementing the fast marching eikonal solver: Spherical versus Cartesian coordinates: *Geophysical Prospecting*, **49**, 165–178.
- Buske, S., and Kästner, U., 2004, Efficient and accurate computation of seismic traveltimes and amplitudes: *Geophysical Prospecting*, **52**, 313–322.
- Chopp, D. L., 2001, Some improvements of the fast marching method: *SIAM Journal of Scientific Computation*, **23**, 230–244.
- 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.
- Grechka, V. Y., and McMechan, G. A., 1996, 3-D two-point ray tracing for heterogeneous, weakly transversely isotropic media: *Geophysics*, **61**, 1883–1894.
- 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*, **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.
- Kim, S., and Cook, R., 1999, 3D traveltimes computation using second-order ENO scheme: *Geophysics*, **64**, 1867–1876.
- Manuel, C. D. and Uren, N. F., 2000, Pre-stack depth migration of seismic multiples: *Exploration Geophysics*, **31**, 328–332.
- Manuel, C. D., Uren, N. F. and Lambert, G. P., 2001, The generalised exploding reflector method: 71st Ann. Internat. Mtg. Soc. Explor. Geophys., Expanded Abstracts, 1167–1170.
- 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., 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 traveltimes tomography of the crust and lithosphere: *Advances in Geophysics*, **46**, 81–198.
- Rawlinson, N., and Kennett, B. L. N., 2004, Rapid estimation of relative and absolute delay times across a network by adaptive stacking: *Geophysical Journal International*, **157**, 332–340.
- Rawlinson, N., and Sambridge, M., 2004, Wave front evolution in strongly heterogeneous layered media using the fast marching method: *Geophysical Journal International*, **156**, 631–647.
- 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.
- Sethian, J. A., 1996, A fast marching level set method for monotonically advancing fronts: *Proceedings of the National Academy of Science*, **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 traveltimes computation using the fast marching method: *Geophysics*, **64**, 516–523.
- Sethian, J. A., 2001, Evolution, implementation, and application of level set and fast marching methods for advancing fronts: *Journal of Computational Physics*, **169**, 503–555.
- Steck, L. K., Thurber, C. H., Fehler, M., Lutter, W. J., Roberts, P. M., Baldrige, W. S., Stafford, D. G., and Sessions, R., 1998, Crust and upper mantle *P* wave velocity structure beneath Valles caldera, New Mexico: Results from the Jemez teleseismic tomography experiment: *Journal of Geophysical Research*, **103**, 24301–24320.
- 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.
- van Trier, J., and Symes, W. W., 1991, Upwind finite-difference calculation of traveltimes: *Geophysics*, **56**, 812–821.
- vanDecar, J. C., James, D. E., and Assumpção, M., 1995, Seismic evidence for a fossil mantle plume beneath South America and implications for plate driving forces: *Nature*, **378**, 25–31.
- Vidale, J. E., 1988, Finite difference calculations of traveltimes: *Bulletin of the Seismological Society of America*, **78**, 2062–2076.
- Zelt, C. A., and Smith, R. B., 1992, Seismic traveltimes inversion for 2-D crustal velocity structure: *Geophysical Journal International*, **108**, 16–34.
- Zhao, P., Uren, N. F., Wenzel, F., Hatherly, P. J. and McDonald, J. A., 1998, Kirchhoff diffraction mapping in media with large velocity contrasts: *Geophysics*, **63**, 2072–2081.