

Wave front evolution in strongly heterogeneous layered media using the fast marching method

N. Rawlinson and M. Sambridge

Research School of Earth Sciences, Australian National University, Canberra ACT 0200, Australia. E-mail: nick@rsees.anu.edu.au

Accepted 2003 October 6. Received 2003 September 19; in original form 2002 October 8

SUMMARY

The fast marching method (FMM) is a grid based numerical scheme for tracking the evolution of monotonically advancing interfaces via finite-difference solution of the eikonal equation. Like many other grid based techniques, FMM is only capable of finding the first-arriving phase in continuous media; however, it distinguishes itself by combining both unconditional stability and rapid computation, making it a truly practical scheme for velocity fields of arbitrary complexity. The aim of this paper is to investigate the potential of FMM for finding later arriving phases in layered media. In particular, we focus on reflections from smooth subhorizontal interfaces that separate regions of continuous velocity variation. The method we adopt for calculating reflected phases involves two stages: the first stage initializes FMM at the source and tracks the incident wave front to all points on the reflector surface; the second stage tracks the reflected wave front by reinitializing FMM from the interface point with minimum traveltime. Layer velocities are described by a regular grid of velocity nodes and layer boundaries are described by a set of interface nodes that may be irregularly distributed. A triangulation routine is used to locally suture interface nodes to neighbouring velocity nodes in order to facilitate the tracking of wave fronts to and from the reflector. A number of synthetic tests are carried out to assess the accuracy, speed and robustness of the new scheme. These include comparisons with analytic solutions and with solutions obtained from a shooting method of ray tracing. The convergence of traveltimes as grid spacing is reduced is also examined. Results from these tests indicate that wave fronts can be accurately tracked with minimal computational effort, even in the presence of complex velocity fields and layer boundaries with high curvature. Incident wave fronts containing gradient discontinuities or shocks also pose no difficulty. Further development of the wave front reinitialization scheme should allow other later arrivals such as multiples to be successfully located.

Key words: fast marching method, finite-difference methods, ray tracing, reflection seismology, seismic-wave propagation, traveltime.

1 INTRODUCTION

Seismic traveltimes are used in a variety of ways to construct images of the Earth's interior. One commonly used approach, called seismic traveltime tomography, constructs 2-D and 3-D maps of wave speed variation by inversion of source–receiver traveltimes. Tomographic methods work by representing the seismic structure of the region being mapped with a set of model parameters, predicting the model data for a given *a priori* model, and then adjusting the values of the model parameters to better match the model data with the observed data. Much has been learned concerning the structure of the Earth at local (e.g. Aki *et al.* 1977; Thurber 1983; Oncescu *et al.* 1984; Hole *et al.* 1992; Eberhart-Phillips & Michael 1993; Darbyshire *et al.* 1998; Graeber & Asch 1999; Zelt *et al.* 2001; Graeber *et al.* 2002), regional (e.g. Spakman 1991; Widiyantoro & van der Hilst 1997; Gorbato *et al.* 2000, 2001) and global (e.g. Dziewonski &

Woodhouse 1987; Inoue *et al.* 1990; Morelli 1993; van der Hilst *et al.* 1997; Bijwaard & Spakman 2000) scales through the application of seismic traveltime tomography. Another method of imaging structure that is frequently used by the exploration industry involves mapping the recorded wavefield directly into model space. This is done using coincident reflection data, which allows reflected phases to be associated with the Earth's impulse response and hence Earth structure. A popular method for producing images from this class of data is Kirchhoff migration (Sena & Toksöz 1993; Gray & May 1994; Epili & McMechan 1996; Buske 1999), which requires the traveltimes to all points in the medium to be calculated.

The important role seismic traveltimes play in modern seismology is reflected in the number of methods that have been devised for their prediction. These include shooting and bending methods of ray tracing, finite-difference solution of the eikonal equation on a grid, wave front construction using local ray tracing and network

methods. Shooting methods of ray tracing formulate the ray equation as an initial-value problem, which allows a complete ray path to be traced provided the source coordinate and initial ray direction are known. The boundary value problem of locating a source–receiver ray path can then be solved by shooting rays through the medium and using information from the computed paths to update the initial ray trajectories. Numerous ray tracing schemes that use this basic approach have been developed (e.g. Julian & Gubbins 1977; Cassell 1982; Červený 1987; Sambridge & Kennett 1990; Bulant 1996; Rawlinson *et al.* 2001a). The bending method of ray tracing iteratively adjusts the geometry of an initial arbitrary path that joins the source and receiver until it becomes a true ray path (i.e. it satisfies Fermat's principle). Bending methods are well suited to media in which velocity varies smoothly (e.g. Julian & Gubbins 1977; Um & Thurber 1987; Prothero *et al.* 1988; Grechka & McMechan 1996) and have also been developed for layered media (e.g. Pereyra *et al.* 1980; Zhao *et al.* 1992).

The principal drawbacks of ray tracing are related to robustness, speed and ray selection. In the presence of velocity variations, both shooting and bending schemes may fail to converge on true two point ray paths; this lack of robustness becomes more evident as the complexity of the medium increases. Ray tracing can also be a time consuming process, especially if a large number of sources and/or receivers are involved and the medium is 3-D. The final difficulty, that of ray selection, is due to the possibility of more than one traveltime and path existing for a source–receiver pair. Usually, one attempts to pick the first-arrival from a seismogram (or the first-arrival of a later phase), but it is difficult to ascertain with both shooting and bending whether the predicted traveltime is a global minimum. Of course, stability, speed and ray selection all trade-off with each other to some extent. For example, increasing the number of initial rays in a shooting scheme may result in more source–receiver paths being found, but will also increase computation time. A comprehensive exploration of ray theory can be found in Červený (2001).

Another more recently developed class of methods for predicting traveltimes seek finite-difference solutions to the eikonal equation throughout a gridded velocity field. Vidale (1988) proposes a finite-difference scheme that progressively integrates traveltimes along an expanding square in 2-D, and later extends the method to 3-D (Vidale 1990). Since then, numerous studies have sought to develop and refine this basic scheme. For example, Hole & Zelt (1995) and Afnimar & Koketsu (2000) improve the stability of the expanding square formalism so that it may be used more effectively in complex media; Qin *et al.* (1992) use the same finite-difference operators as Vidale (1988) but calculates the traveltime field along the shape of the expanding wave front; Podvin & Lecomte (1991) and van Trier & Symes (1991) seek to improve algorithm stability by using Huygen's principle in the finite-difference approximation and upwind difference operators, respectively.

A characteristic shared by most finite-difference schemes is that they are only capable of finding the absolute first-arrival traveltimes in continuous media. While this may be advantageous for some tomographic schemes and Kirchhoff migration, there are many occasions in which later arrivals are of interest. However, certain classes of later arriving phases such as reflections from specified interfaces can be calculated with finite-difference schemes (Podvin & Lecomte 1991; Riahi & Juhlin 1994; Hole & Zelt 1995). Another limitation of finite-difference schemes is that they may suffer from stability problems; in particular, an expanding square formalism may breach causality in the presence of large velocity gradients (Qin *et al.* 1992).

Shortest path ray tracing (SPR) offers an alternative grid method for calculating traveltimes. A grid of nodes is specified within the velocity medium and a network or graph is formed by connecting neighbouring nodes with traveltime path segments (Nakanishi & Yamaguchi 1986; Moser 1991; Fischer & Lees 1993; Cheng & House 1996). The problem then reduces to one of finding the path of minimum traveltime between the source and the receiver through the network, which may be solved using Dijkstra-like algorithms. Similar to finite-difference solutions of the eikonal equation, SPR is only capable of finding first-arrivals unless specific conditions, such as reflections from an interface, are imposed (Moser 1991). Like any grid based method, the accuracy of SPR will improve as node spacing decreases, but having large angles between adjacent connectors can result in a poor approximation to the true path. Finite-difference methods generally offer a better compromise between speed and accuracy than SPR methods (Leidenfrost *et al.* 1999), but the latter is often considered to be more numerically stable (e.g. Cheng & House 1996).

Wave front construction (WC) is another recently developed numerical technique for calculating traveltimes (Vinje *et al.* 1993; Ettrich & Gajewski 1996; Lambaré *et al.* 1996; Lucio *et al.* 1996; Vinje *et al.* 1999). The basic principle WC employs is that discrete propagation of a wave front through a medium can be achieved by using local ray tracing from each wave front surface. Rays are traced for a given time step, with the end points of rays describing the geometry of the new wave front. If diverging rays cause parts of the wave front to be poorly sampled, new rays are introduced by interpolation. Advantages of WC compared with most grid based methods include the ability to construct multi-valued traveltime fields and the accurate calculation of amplitudes. However, WC tends to be slower than grid-based methods for the same level of traveltime accuracy (see Leidenfrost *et al.* 1999) and portions of the wave front with high curvature can be poorly approximated. Lambaré *et al.* (1996) introduce a uniform ray density criterion in phase space to help overcome this problem.

Gudmundsson *et al.* (1994) devise a method of wave front construction on a regular grid, which uses ray perturbation theory to estimate traveltimes to nodes in the neighbourhood of the new wave front location and then extracts the shape of the wave front using a contouring routine. Although later arrivals are not found with this scheme, undersampling of the wave front is avoided. A comparison of finite-difference, SPR and WC methods can be found in Leidenfrost *et al.* (1999).

The problem of tracking the evolution of a monotonically advancing interface is not restricted to seismic wave fronts; there are many other areas of science in which this problem needs to be solved. A recently introduced technique called the fast marching method or FMM (Sethian 1996; Sethian & Popovici 1999), was developed with this in mind. To date, FMM has been applied to a wide variety of problems including photolithographic development, geodesics, deposition of sediments, medical imaging and optimal path planning (Sethian 1996, 1999, 2001). At its core, FMM is a grid based numerical algorithm, which tracks the evolving interface along a narrow band of nodes where the values are updated by solving the eikonal equation using upwind entropy satisfying finite-difference approximations to the gradient term. A particular strength of FMM is that it combines unconditional stability with rapid computation, a property that has not been demonstrated for any of the schemes discussed previously. The unconditional stability of FMM comes from properly addressing the development and propagation of gradient discontinuities in the evolving wave front (Sethian 1996; Sethian & Popovici 1999). Another useful property of FMM is that the order of




Figure 1. A schematic illustration of how the upwind scheme obtains the entropy satisfying solution. (a) Propagating wave front with a gradient discontinuity results in a swallow tail; (b) the viscous-limit solution is the first-arrival wave front; (c) the entropy-satisfying upwind scheme (eq. 3) finds the first-arrival to the point T_{ij} by separately considering the propagation of the wave front on either side of the gradient discontinuity (see text for more details). Black grid nodes have known traveltimes while grey grid nodes have no associated traveltimes.

accuracy of the method can be controlled by using the appropriate upwind finite-difference operators. In the seismic context, Sethian & Popovici (1999) introduce a first-order FMM for calculating traveltimes in 3-D media, Alkhalifah & Fomel (2001) compare first-order FMMs in Cartesian and spherical coordinates, and Popovici & Sethian (2002) introduce a higher-order scheme and apply it to imaging of reflection data.

In this paper, we apply FMM to layered media with the aim of tracking later phase-specific arrivals. In particular, we focus on reflected phases from layer boundaries. Reflection traveltimes are of great importance in many applications including reflection and wide-angle tomography; a fast unconditionally stable scheme for calculating these arrivals would be of significant value. The approach we adopt is to apply a local triangulation scheme to associate irregularly spaced boundary nodes with adjacent nodes from the regular velocity grid. First-, second- and higher-order solutions of the eikonal equation may be simply applied within the regular grid, while wave fronts in the vicinity of the layer boundary are solved using a first-order version of FMM for triangulated meshes. Reflected waves may be tracked by using a two-stage scheme, which involves propagating the incident wave front to every point on the interface, and then reinitializing the scheme from the interface node with minimum traveltime in order to track the reflected wave front back through the medium.

After presenting the necessary theory, we make a brief examination of the comparative accuracy of first-, second- and third-order FMM schemes in order to identify the scheme that offers the optimum balance between accuracy and computation time. Two examples demonstrating the unconditional stability of the scheme in highly heterogeneous media are also presented. Examples of FMM solutions in layered media are then given, focusing on the accuracy of reflected traveltimes for smoothly varying interfaces. This is achieved by comparing FMM results with those produced by a very accurate ray-shooting method developed by Rawlinson *et al.* (2001a). Our final results demonstrate the stability of the scheme in the presence of highly complex layered structures.

2 METHOD

2.1 FMM in continuous media

We begin by giving a brief description of FMM in continuous media. For more details, refer to Sethian (1996), Sethian & Popovici (1999), Sethian (1999) and Popovici & Sethian (2002). The eikonal equation states that the magnitude of the traveltime gradient at any point along a wave front is equal to the inverse of the velocity at that point and

may be written as

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

where $\nabla_{\mathbf{x}}$ is the gradient operator, T is traveltime and $s(\mathbf{x})$ is slowness.

A significant obstacle for finite-difference methods that seek to solve the eikonal equation for the first-arrival traveltime field is that the wave front may be discontinuous in gradient. These discontinuities are especially apparent in velocity media that cause multipathing; the wave front self-intersects but the later arriving information is discarded (e.g. Fig. 1). The eikonal equation cannot be easily solved in the presence of gradient discontinuities because the equation itself requires $\nabla_{\mathbf{x}} T$ to be defined. One way of dealing with this problem is to seek ‘weak solutions’, which result in a continuous $T(\mathbf{x})$ but not necessarily a continuous $\nabla_{\mathbf{x}} T$. One way of obtaining a weak solution is to solve the ‘viscous’ version of the eikonal equation:

$$|\nabla_{\mathbf{x}} T| = s(\mathbf{x}) + \epsilon \nabla_{\mathbf{x}}^2 T \quad (2)$$

as $\epsilon \rightarrow 0$ (the viscous limit), where the parameter ϵ controls the smoothness imposed on the solution. The limit of smooth solutions is a weak solution that corresponds to the first-arriving wave front. By definition, the propagation of a first-arriving wave front satisfies an entropy condition because it can only pass through a point once; hence, information can only be lost or retained as the wave front evolves. Sethian & Popovici (1999) state this entropy condition as ‘once a point burns, it stays burnt’. The unconditional stability of FMM comes from strict enforcement of this entropy condition (see Sethian 1996; Sethian & Popovici 1999, for more details).

Entropy satisfying weak solutions may be obtained by using upwind gradient operators in eq. (1), which take into account the direction of flow of information. Nodes that reside downwind from the wave front should only be updated by using traveltime values from nodes that lie upwind. The upwind scheme we employ is one that has been used by a number of authors including Sethian & Popovici (1999), Chopp (2001) and Popovici & Sethian (2002) and may be expressed:

$$\left[\begin{array}{l} \max(D_a^{-x} T, -D_b^{+x} T, 0)^2 \\ + \max(D_c^{-y} T, -D_d^{+y} T, 0)^2 \\ + \max(D_e^{-z} T, -D_f^{+z} T, 0)^2 \end{array} \right]_{ijk}^{1/2} = s_{i,j,k}, \quad (3)$$

where T is traveltime, (i, j, k) are 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.

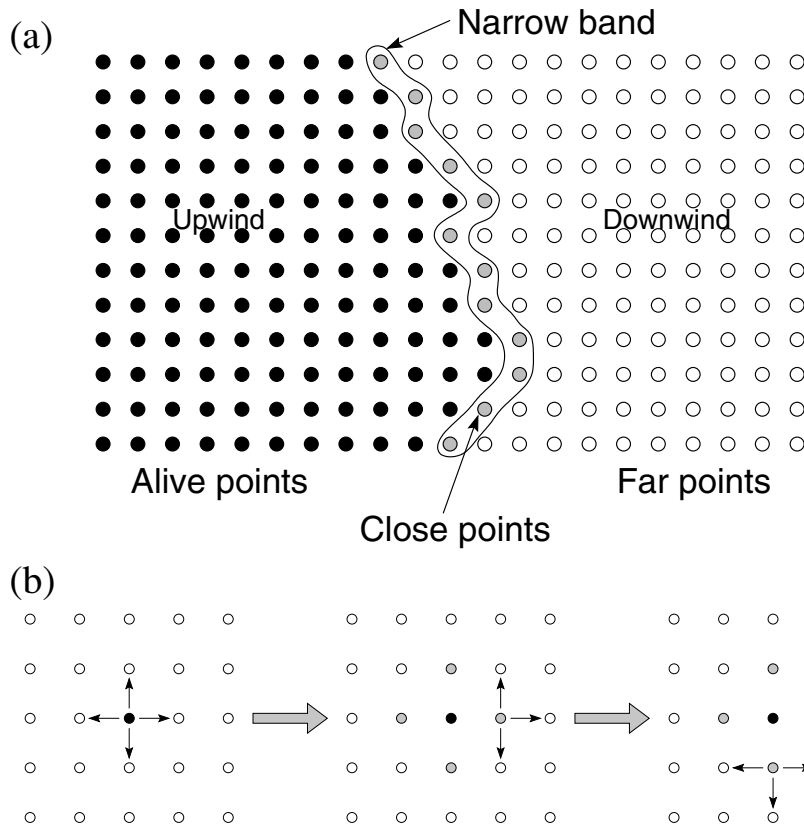


Figure 2. (a) The principle of the narrow-band method. (b) Example of how the narrow band evolves from a source point. See the text for more details.

For example, the first three upwind operators for $D^{-x}T_i$ are:

$$\begin{aligned}
 D_1^{-x}T_i &= \frac{T_i - T_{i-1}}{\delta x} \\
 D_2^{-x}T_i &= \frac{3T_i - 4T_{i-1} + T_{i-2}}{2\delta x} \\
 D_3^{-x}T_i &= \frac{11T_i - 18T_{i-1} + 9T_{i-2} - 2T_{i-3}}{6\delta x},
 \end{aligned}
 \tag{4}$$

where δx is the grid spacing in x . These operators are easily derived by appropriate summation of the Taylor series expansions for T_{i-1} , T_{i-2} and T_{i-3} . Which operator is used in eq. (3) depends on the availability of upwind traveltimes and the maximum order allowed. First-order accurate schemes only use D_1 operators, second-order accurate schemes preferentially use D_2 operators and third-order schemes preferentially use D_3 operators. Strictly speaking, the second- and third-order methods are really mixed order because they will use lower-order approximations when causality does not permit the use of the required operator. For example, if we implement a third-order method and $T_{i-1} > T_{i-2} > T_{i-3}$, then the operator used would be $D_3^{-x}T_i$; if $T_{i-1} > T_{i-2} < T_{i-3}$, then we would use $D_2^{-x}T_i$; while if $T_{i-1} < T_{i-2}$ we would have to resort to $D^{-x}_1T_i$. For the remainder of the paper, ‘ n th-order scheme’ implies that n is the maximum order operator allowed, and not that the scheme is strictly n th-order accurate.

A simple example of the entropy-satisfying nature of eq. (3) is shown in the schematic diagram of Fig. 1. Here, a wave front containing a gradient discontinuity is propagated. The propagation of the complete front produces a swallow tail solution (Fig. 1a). Using a non-zero viscosity term (eq. 2) will smooth out the discontinuity and eliminate the need for later arrivals. The viscous limit solution (Fig. 1b) also discards the later-arriving swallow tail, but preserves

the discontinuity and hence correctly calculates the first-arrival wave front. The entropy-satisfying upwind scheme will also produce the first-arrival wave front by separately considering information on either side of the gradient discontinuity (Fig. 1c). In quadrant A, eq. (3) with first-order operators reduces to the quadratic:

$$\left(\frac{T_{i,j} - T_{i-1,j}}{\delta x}\right)^2 + \left(\frac{T_{i,j} - T_{i,j-1}}{\delta z}\right)^2 = (s_{i,j})^2.
 \tag{5}$$

Of the two possible solutions, the larger one corresponds to the wave front segment in quadrant A impinging upon (i, j) . Similarly in quadrant B, the appropriate quadratic is

$$\left(\frac{T_{i+1,j} - T_{i,j}}{\delta x}\right)^2 + \left(\frac{T_{i,j} - T_{i,j-1}}{\delta z}\right)^2 = (s_{i,j})^2
 \tag{6}$$

and again the larger solution corresponds to the wave segment in quadrant B impinging upon (i, j) . Of the two options from A and B, the correct solution is the one with minimum traveltime. In Fig. 1(c), both traveltimes are the same and (i, j) becomes a point of gradient discontinuity, but this need not be the case in general.

The upwind scheme given by eq. (3) describes how to calculate new traveltimes using known traveltimes from adjacent gridpoints. 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 values of T to larger values of T . To achieve this, FMM systematically constructs traveltimes T in a downwind fashion from known values upwind by employing a *narrow-band* approach. The narrow-band concept is illustrated in Fig. 2(a); *alive* points have their values correctly calculated, *close* points lie within the narrow band and have trial values, and *far* points have no values calculated. Trial values are calculated using eq. (3)




Figure 3. The principle behind the determination of reflection traveltimes. The incident wave front is tracked to all points on the interface. All velocity nodes are then set to *far*, all interface nodes are set to *close* and the scheme is re-initialized to track the reflected wave front.

with *alive* points only, which by definition lie upwind of the *close* points. The shape of the narrow-band approximates the shape of the first-arrival wave front, and the idea is to propagate the band through the grid until all points become *alive*.

The narrow band is evolved by identifying the *close* point with minimum traveltime and tagging it as an *alive* point. All neighbours of this point that are not *alive* are tagged as *close* and have their traveltimes computed (or recomputed) using eq. (3). Repetition of this process allows the traveltime to all gridpoints to be calculated. Fig. 2(b) illustrates several evolution steps of the narrow-band method from a source point. Choosing the *close* points with minimum traveltime guarantees that *alive* points are not evaluated with partial information. In general, adjacent *close* points must lie upwind or downwind of each other; the *close* point with minimum traveltime will always lie upwind of adjacent *close* points. A heap sort algorithm is used to rapidly locate the gridpoint in the narrow band with minimum traveltime. Since information in a heap sort is stored in a binary tree, FMM will have an operation count of $O(M \log M)$ where M is the total number of gridpoints.

Although ray paths are not explicitly found by FMM, they may be calculated *a posteriori* using the computed traveltime field. The easiest way to do this is to start at each receiver and follow the traveltime gradient ∇T back to the source. ∇T can be approximated by finite differences in a manner similar to the upwind scheme. The rays traced through the models in Section 3 are computed using this approach.

2.2 FMM in layered media

We adapt the FMM scheme described above to allow phases other than the absolute first-arrival to be tracked in layered media. In particular, we focus on reflections from subhorizontal interfaces. The two principal difficulties involved in the introduction of reflections are as follows. (1) The entropy condition that is central to solving the eikonal equation requires that once the first-arrival traveltime to a point has been found, the wave front cannot pass through that point again. To permit reflected waves, this condition must be breached, but not in a way that undermines the correct calculation of traveltimes. (2) Layer boundaries that vary with depth must be described by an irregular distribution of interface nodes, which do not conform with the regular velocity grid.

A relatively simple way of dealing with the problem of a multi-valued traveltime field is to adopt the following two stage approach. First, track the incident wave front to all points on the interface. Then set all velocity nodes above the interface to *far* and all interface nodes to *close* and reinitialize the FMM scheme to track the reflected wave front. Although the narrow band at the start of the second stage ‘surrounds’ the interface and not the wave front, causality is not breached because the first-arrival reflection wave front only impinges on the interface once at every point. Fig. 3 illustrates the principle behind our two stage scheme.

The shape of a reflected wave front has a strong dependence on the geometry of the reflecting interface, so it is essential that layer boundaries are accurately represented by the computational grid. The scheme we adopt defines a set of interface nodes by locating every intersection point between the interface and the cell boundaries of the rectangular velocity grid (see Fig. 4). This results in an irregular distribution of interface nodes embedded within a regular grid of velocity nodes. Traveltime information flowing to and from the interface nodes cannot be calculated using eq. (3) because there is no rectangular grid association between neighbouring nodes. Instead, we use an irregular mesh of triangles (for 2-D velocity media only) to locally suture the interface nodes to adjacent velocity nodes (Fig. 4). The meshing routine is designed to minimize the presence of obtuse triangles. Traveltimes are updated across the irregular mesh by using the triangular first-order upwind scheme of Sethian (1999), which is outlined below. Advantages of using irregular grids locally at an interface, rather than globally throughout the medium, include reduced computational requirements and the ability to retain higher-order accuracy within the rectangular mesh.

Consider the interface node O shown in Fig. 5(a), which is surrounded by a mesh of four triangles. Like FMM for a rectangular grid, the traveltime prediction for each of these triangles needs to be calculated and the minimum chosen as the correct solution. If we select one of the four triangles and assume that the traveltimes to points A (T_A) and B (T_B) are known (see Fig. 5b), then the eikonal equation can be used to determine the traveltime to point O (T_O) with first-order accuracy by assuming that the traveltime gradient within the triangle is constant. If t is the local traveltime from T_A to T_O , and $T_B > T_A$, then a quadratic equation for t is given by (see




Figure 4. Interface nodes are locally sutured to adjacent velocity nodes via a triangular mesh. Thick dashed lines represent irregular mesh boundaries and the solid black curve represents the true layer boundary. Velocity and interface nodes are denoted by grey circles and triangles, respectively.




Figure 5. Updating traveltimes in a triangulated domain. (a) In order to update point O, traveltimes calculations need to be considered for all four triangles. (b) Using known traveltimes at point A and B to update point O. The traveltimes gradient is assumed to be constant within a triangle. Velocity and interface nodes are denoted by grey circles and triangles, respectively.




Figure 6. Illustration of why acute triangulation is necessary to preserve first-order accuracy in a triangulated domain. (a) A plane wave front impinging on a network of acute triangles; (b) on right angled triangles; (c) on obtuse triangles. In the latter case, it is possible for wave fronts to pass through only one of the bounding nodes (A, B or C) before arriving at O. This leads to traveltimes that are not first-order accurate.

Sethian 1999 for a derivation):

$$(a^2 + b^2 - 2\mathbf{a} \cdot \mathbf{b})t^2 + 2u(\mathbf{a} \cdot \mathbf{b} - b^2)t + (b^2u^2 - s_0^2[a^2b^2 - (\mathbf{a} \cdot \mathbf{b})^2]) = 0, \tag{7}$$

where s_0 is slowness at O, $u = T_B - T_A$, \mathbf{a} and \mathbf{b} are displacement vectors (see Fig. 5b), and $a = |\mathbf{a}|$ and $b = |\mathbf{b}|$.

In order to correctly use eq. (7), the traveltimes gradient vector at O must lie between points A and B (see Fig. 5b) and $u < t$. It can be shown that the traveltimes gradient condition is satisfied provided (see Sethian 1999):

$$\frac{\mathbf{a} \cdot \mathbf{b}}{b} < \frac{b(t - u)}{t} < \frac{a^2b}{\mathbf{a} \cdot \mathbf{b}}. \tag{8}$$

If both conditions are satisfied, then the solution is $T_O = t + T_A$; otherwise the solution is $T_O = \min\{bs_0 + T_A, as_0 + T_B\}$. In the example shown in Fig. 5(a), up to four values for T_O (one for each triangle) may be obtained, but the minimum value is always selected. Note that the expressions given in eqs (7) and (8) avoid the trigonometric functions used in the equivalent expressions given by Sethian (1999).

An important restriction in applying FMM to triangulated domains is that angles between adjacent mesh lines emanating from a node must be acute. To understand why this restriction exists, consider the sequence of three diagrams in Fig. 6. In all cases, we are aiming to calculate the traveltimes to point O from an impinging plane wave front. Mesh lines emanate from O at acute angles (72°) in Fig. 6(a), and it is clear that a wave front impinging on O from any angle must first pass through at least two adjacent nodes, in which case the update will be first-order accurate since eq. (7) will be used. In Fig. 6(b), mesh lines emanate at right angles, and the wave front must pass through two adjacent points before reaching O except when the wave front is perpendicular to AO, BO, CO or DO.

Fig. 6(b) shows a wave front perpendicular to AO, with traveltimes given by $T_O = T_A + s_0AO$, which is also first-order accurate.

Fig. 6(c) illustrates the case of obtuse triangles, with three grid lines emanating from O with angular separations of 120° . If we consider the triangle AOB, it is clear that wave fronts with normals that lie in the zone defined by θ will pass through both A and B before reaching O. However, wave fronts with normals that lie within either ϕ_1 or ϕ_2 will only pass through A or B, respectively, before reaching O. For the wave front in Fig. 6(c), the traveltimes to O will be approximated by $T_O = T_A + s_0AO$, which is not first-order accurate. In general, the percentage of all incidence angles that result in less than first-order accurate estimates of traveltimes is given by

$$r = 100 \left(\frac{2\psi - \pi}{\psi} \right), \tag{9}$$

where ψ is the angle that subtends AOB and $\pi/2 \leq \psi \leq \pi$. Ultimately, the use of obtuse triangles will reduce the accuracy of the scheme because it will no longer be truly first-order accurate.

For the local adaptive triangulation scheme illustrated in Fig. 4, it is possible for obtuse triangles to be generated. This can only occur when the interface cuts a corner off a cell in the manner illustrated in Fig. 7. To fix this problem, we can simply use a different splitting strategy that involves replacing the mesh line AC with BO; however, this is only required when the point O is being updated. For all other updates (i.e. to A, B, C or D), the original strategy is used. The use of both strategies ensures that all nodes have mesh lines emanating with acute angular separation when they are updated. This guarantees that the correct plane-wave approximation is used to update every point within the triangulated domain.

So far, we have only considered reflected phases in layered media, but an incident wave front will also transmit through a layer boundary. One way of calculating transmissions is to apply the triangulation routine to both sides of an interface and use a one stage




Figure 7. Dealing with obtuse angles by using a different splitting strategy.

FMM. However, transmitted rays are generally less deviated by an interface than reflected rays, so it is possible to obtain suitably accurate refraction traveltimes without explicitly including interfaces in the parametrization. Thus, the velocity nodes on either side of an interface have correct layer velocities, but interface nodes are not included. In this approach, the interface is effectively approximated by a sharp velocity gradient, which makes computation a simpler task and allows structures such as pinch outs to be more easily represented.

3 RESULTS

The class of structure that our FMM scheme is applied to consists of 2-D subhorizontal layers separated by smooth interfaces. In principle, the irregular local mesh scheme that associates velocity nodes with adjacent interface nodes could be extended to 3-D using a tetrahedral mesh rather than a triangular one. However, as we wish to demonstrate the stability, accuracy and further potential of our scheme to find other later arrivals, examples in 2-D are sufficient. We also supply the CPU times required to execute most of our FMM examples. All calculations were performed in double precision on a Sun Ultra 5; CPU times include reading the velocity grid from disc.

Application of the FMM scheme requires the velocity field to be discretely sampled; the node spacing chosen, along with the order of the upwind difference operator used, will determine the accuracy of the calculated traveltimes. Thus, the velocity field and interface surface should both be described in a way that makes discrete sampling at any specified grid spacing a straightforward task. Another important consideration when higher-order operators are used is the smoothness of the velocity continuum. True second-order accuracy requires the velocity field to be continuous in gradient (C_1 continuity), while true third-order accuracy requires the velocity field to be continuous in curvature (C_2 continuity).

With these considerations in mind, we use cubic B-spline functions in parametric form (see Bartels *et al.* 1987) to describe the underlying velocity field and interface structure. Variations in interface and velocity structure are defined by a separate set of interface and velocity vertices, which control the shape of the cubic functions. Discrete sampling of these functions to any desired resolution is easily accomplished. Cubic B-splines have many useful properties including local control and C_2 continuity, the latter of which is crucial for our implementation of higher-order finite-difference operators to solve the eikonal equation. Smoothness and local control also make them an attractive parametrization in seismic tomography (e.g. Rawlinson *et al.* 2001a), which is one possible application for

FMM. The examples presented below examine the performance of FMM in the presence of crustal scale structure, but the results are also applicable to structures at different scales.

3.1 Examples in continuous media

The first set of results we present analyses the accuracy and robustness of FMM in continuous media. These results will be of benefit in helping to understand the behaviour of FMM in layered media, examples of which are presented in the next section. A constant velocity medium is a good starting point for analysing the performance of FMM because analytic solutions are easy to calculate. Fig. 8(a) shows wave fronts and rays emanating from a point source in a medium of velocity 6.0 km s^{-1} . Rays are traced to a set of 21 receivers positioned 5 km apart on the upper surface of the medium. Traveltimes to these 21 receivers form the basis of our analysis. The rms traveltimes T_{rms} of the exact solution is provided to help facilitate error comparison between this and subsequent examples.

The accuracy of the first-order FMM for different grid sizes is illustrated in Fig. 8(b). The misfit between the numerical solution and the analytic solution is defined by $\Delta t = t_{\text{fmm}} - t_a$ where t_{fmm} is the solution produced by FMM and t_a is the analytic (exact) solution. The four computational grids that are used for FMM have node spacings of 1000, 500, 250 and 125 m. It is easy to see from Fig. 8(b) that the accuracy of FMM increases with decreasing node spacing, as one would expect. In fact, the rms traveltimes misfit (see Table 1) is reduced by approximately 41 per cent each time the grid spacing is halved. The traveltimes error is zero at $x = 0$ because the finite-difference solution is exact along horizontal and vertical grid lines. As expected, traveltimes error is maximized when rays are oriented near 45° to the x and z axes.

When second-order difference operators are used, the accuracy of FMM is greatly improved (Fig. 8c). When computational cost is considered (Table 1), it is clear that the first-order scheme can only achieve similar accuracy to the second-order scheme with an order of magnitude increase in computational effort. For example, the rms error of the first-order scheme with a grid spacing of 125 m is 32.7 ms and computation time is 5.3 s; the rms error of the second-order scheme with a grid spacing of 1000 m is 35.5 ms and computation time is only 0.1 s. At first glance, these results may suggest that implementing a third-order scheme will result in similar improvements over the second-order scheme. However, Fig. 8(d) clearly shows that this is not the case. In fact, the third-order scheme is *worse* than the second-order scheme for all four grid spacings tested.

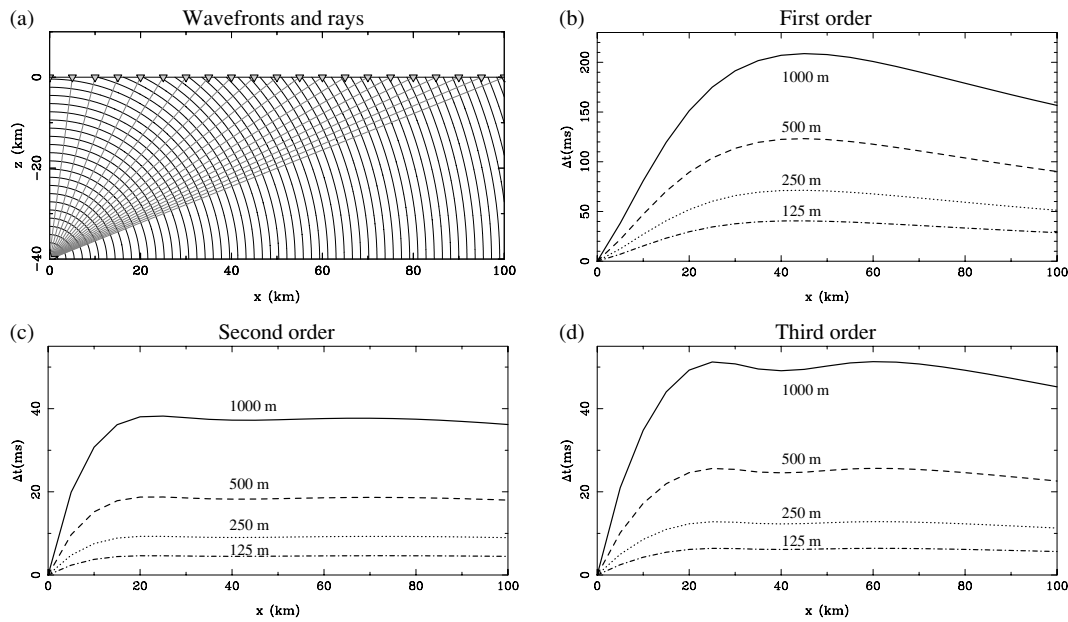


Figure 8. Accuracy of FMM in a uniform velocity medium (6.0 km s^{-1}). (a) Wave fronts (at 0.3 s intervals) are denoted by thin black lines; rays are denoted by thick grey lines; receivers are denoted by grey triangles. (b)–(d) show plots of the difference between numerical and analytic solutions (Δt) using three different FMM schemes with four different grid spacings: 1000, 500, 250 and 125 m. For the exact solution, $T_{\text{rms}} = 9.74 \text{ s}$.

Table 1. Summary of results from tests in a constant velocity medium (see Fig. 8).

Grid spacing (m)	Number of gridpoints	rms error (ms)			CPU time (s)		
		first-order	second-order	third-order	first-order	second-order	third-order
1000	4141	171.1	35.5	46.4	0.1	0.1	0.1
500	16 281	100.3	17.5	23.2	0.3	0.3	0.3
250	64 561	57.7	8.7	11.6	1.2	1.3	1.4
125	257 121	32.7	4.3	5.8	5.3	5.8	6.3

The reason for the failure of the third-order scheme can be traced back to the calculation of the traveltimes in the vicinity of the source point. A wave front produced by a point source will have a very high curvature in the vicinity of the source but will be under-sampled and hence poorly represented by the computational grid. In addition, only a first-order scheme can be used in the neighbourhood of the source. If we apply FMM in a constant velocity medium as before in the Fig. 8 examples, but use analytic solutions in a specified region about the source, we can observe the effects of reducing source error. Fig. 9 compares the error of a second- and a third-order scheme, which use analytic solutions out to a distance of 10 km from the source; beyond this distance, the solutions revert to those given by FMM on a 1000 m grid. The increase in accuracy over the equivalent plots in Figs 8(c) and (d) is dramatic, and clearly indicates that most of the inaccuracy in the FMM traveltimes field is generated in the source region.

The other interesting observation that can be made concerning Fig. 9 is that the third-order scheme (rms error of 0.7 ms) is much more accurate than the second-order scheme (rms error of 5.4 ms). However, the size of this improvement is small compared with the size of the error generated in the source region, which explains why there is no real benefit in using schemes higher than second-order without accounting for near source errors. The reason why the second-order method gives better results than the third-order method in Fig. 8 is due to the traveltimes being overestimated about the source; the second-order scheme tends to underestimate the traveltimes away from the source while the third-order scheme tends to




Figure 9. Comparison of error between second- and third-order schemes in a constant velocity medium (see Fig. 8a) using analytic solutions within a 10 km radius of the source.

overestimate the traveltimes (Fig. 9). In the latter case, the errors are compounded, and in the former case, the errors partially cancel each other out. While this may not always occur, better results cannot necessarily be assumed by using schemes higher than second order.

In realistic applications of FMM, the velocity field is unlikely to be constant, so we also examine the performance of FMM in a medium in which velocity varies linearly with depth. Analytic solutions describe ray trajectories as circular arcs in this case. Fig. 10(a) shows wave fronts and rays from a surface source propagating through a medium with a vertical velocity gradient of 0.1 s^{-1} and a velocity

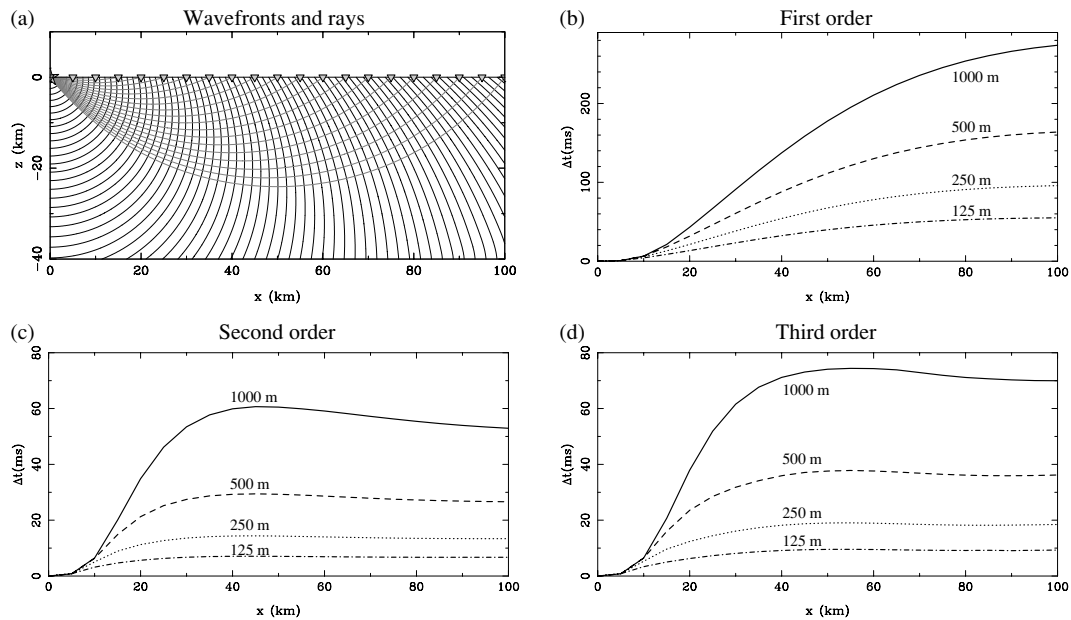


Figure 10. Accuracy of FMM in a uniform velocity gradient medium (vertical gradient of 0.1 s^{-1} with a velocity of 4.0 km s^{-1} at the surface). Information is plotted in the same way as Fig. 8. For the exact solution, $T_{\text{rms}} = 12.98 \text{ s}$.

Table 2. Summary of results from tests in a medium in which velocity varies linearly with depth (see Fig. 10).

Grid spacing (m)	Number of gridpoints	rms error (ms)			CPU time (s)		
		first-order	second-order	third-order	first-order	second-order	third-order
1000	4141	183.1	50.0	62.0	0.1	0.1	0.1
500	16 281	112.1	25.0	31.9	0.3	0.3	0.3
250	64 561	66.8	12.5	16.2	1.2	1.3	1.4
125	257 121	39.0	6.2	8.2	5.3	5.8	6.2

at the surface of 4.0 km s^{-1} . Figs 10(b)–(d) compare the results of using first-, second- and third-order schemes, respectively, with the same grid separations as in Fig. 8. The second-order scheme is again the most accurate, and computation time (see Table 2) is very similar to the constant velocity case. These results suggest that FMM can work well in media which contain significant velocity variations. Since we have demonstrated that no gain in accuracy can be assumed by using third or higher-order schemes, we use a second-order scheme for all remaining examples.

Although we would like to assess the accuracy of FMM in arbitrarily complex media, the lack of analytic solutions in all but a few cases makes this a difficult task. However, the unconditional stability of FMM means that as the grid sized is reduced, the computed solution approaches the true solution. Fig. 11 demonstrates the effects on traveltimes of reducing grid spacing in two velocity models with extreme velocity contrasts. In Fig. 11(a), velocity varies between approximately 1.2 and 8.0 km s^{-1} ; admittedly such variations are unlikely to be encountered in the Earth, but our objective here is to demonstrate the robustness of the scheme. The traveltimes residual plot in Fig. 11(a) is constructed in a similar way to those in Figs 8 and 10 except that the true solution is replaced by a second-order FMM solution with 50 m grid spacing. Successively halving the grid spacing from 1000 down to 125 m results in a convergence towards the 50 m grid solution. In fact, the difference between the 125 m grid solution and the 50 m grid solution is so small that we can treat Δt as an accurate measure of the true error. Thus, the magnitude of the errors we observe are consistent with those determined for the uniform velocity and velocity gradient examples. This is despite

the complexity of the evolving wave front, which includes several propagating wave front gradient discontinuities.

To reinforce the message of unconditional stability and accuracy, we demonstrate the behaviour of FMM in a truly pathological (from the point of view of seismic structure) velocity environment (Fig. 11b). In this example, velocity varies between 1 and 70 km s^{-1} and the velocity structure consists of a fast velocity path winding its way almost sinusoidally through vertically oriented slabs of slow material. The first-arriving wave fronts and rays follow this fast path, resulting in a highly contorted wavefield. However, as the traveltimes residual plot shows (Fig. 11b), traveltimes converge towards the 50 m grid solution, and the difference between traveltimes on a 125 m grid and a 50 m grid is extremely small. The magnitude of the residuals in this example are smaller than those in the Fig. 11(a) example, but it should be noted that the traveltimes in Fig. 11(b) are on average 59 per cent smaller. The CPU times for these examples are not significantly different to the CPU times of the earlier examples, which suggests that computational effort is not dependent on structural complexity, a result we would expect.

3.2 Examples in layered media

We now present a series of examples that illustrate the effectiveness of the new FMM scheme for calculating reflected phases in layered media. Rather than use relatively simple analytic solutions to verify the accuracy of the scheme, we use the ray tracing method of Rawlinson *et al.* (2001a), which can be applied to layered media

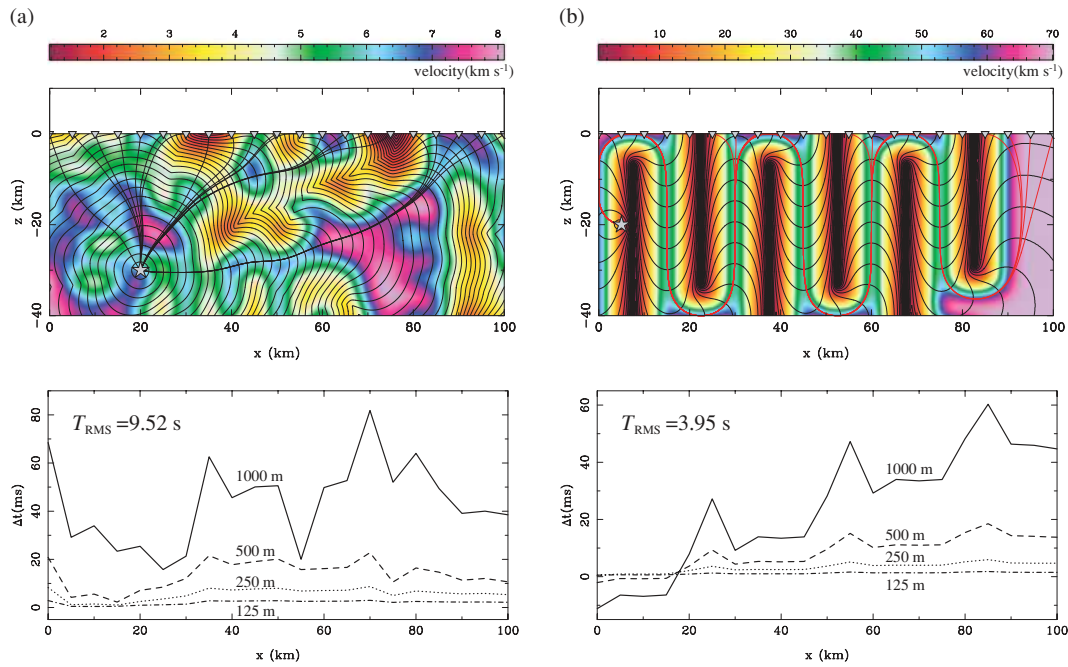


Figure 11. Analysing the stability and accuracy of FMM in complex media. (a) Velocity medium with contrasts as large as 800 per cent. Wave fronts (at 0.3 s intervals) are denoted by thin black lines; rays are denoted by thick black lines. (b) Pathological medium with contrasts as large as 7000 per cent. Wave fronts (at 0.1 s intervals) are denoted by thin black lines; rays are denoted by thick red lines. Sources are denoted by stars and receivers by triangles in both plots. Traveltime error estimates are based on using FMM traveltimes calculated on a 50 m grid.

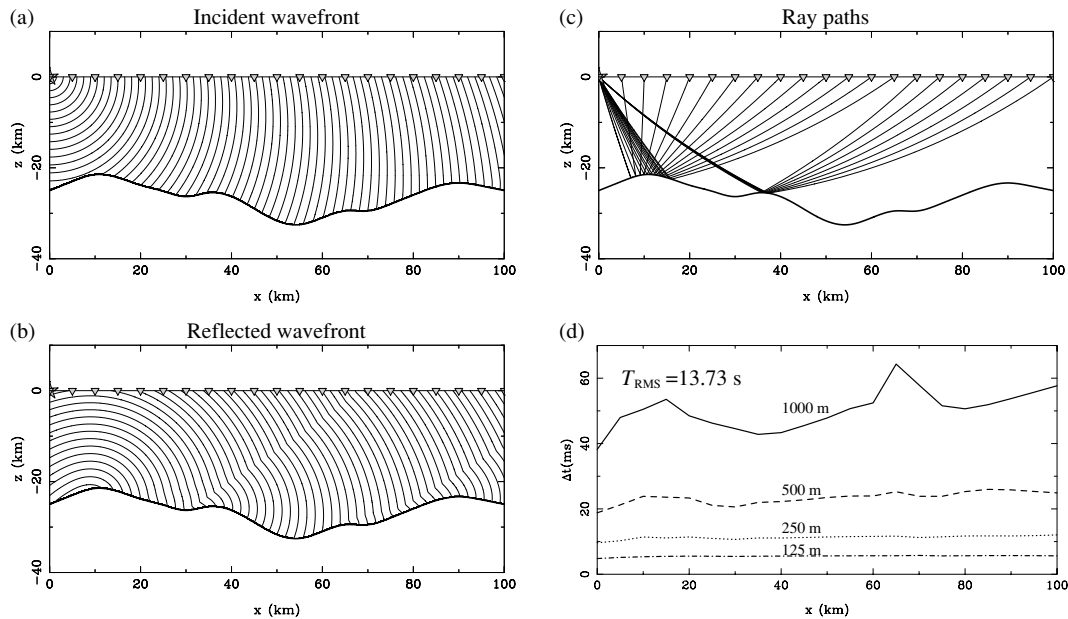


Figure 12. Comparison between the second-order FMM scheme and the shooting method of ray tracing for a reflected phase in a single layer. (a) Incident wave front; (b) reflected wave front; (c) source–receiver ray paths; (d) traveltime residual for four different grid sizes. Wave fronts are contoured at a 0.3 s interval, receivers are denoted by grey triangles and the source by a grey star. The velocity at the surface is 5.0 km s^{-1} and the velocity gradient in the upper layer is 0.04 s^{-1} . The source is located at $(x, z) = (0, 0)$.

containing smoothly varying subhorizontal interfaces. The properties of the ray tracing scheme that are important in this application are: (1) two point rays are located using a shooting method; (2) refracted and reflected rays and traveltimes in layered media can be found; (3) within a layer velocity varies linearly with depth, which permits analytic solutions; (4) layer boundaries are described by cu-

bic B-spline functions; (5) the first-arrivals of specified phases are found. Properties (3) and (4) allow direct comparison of one class of FMM solutions (i.e. no lateral velocity variation within a layer) with the ray tracing solutions.

When a ray is traced using the method of Rawlinson *et al.* (2001a), the component of the initial value scheme that is not analytic

Table 3. Summary of results that compare the accuracy of the second-order FMM with a shooting method of ray tracing in layered media.

Grid spacing (m)	Number of gridpoints	rms error (ms)				CPU time (s)			
		Fig. 12	Fig. 13	Fig. 14	Fig. 15	Fig. 12	Fig. 13	Fig. 14	Fig. 15
1000	4141	50.6	48.6	56.6	39.4	0.1	0.1	0.1	0.1
500	16 281	23.5	23.3	25.1	22.5	0.4	0.4	0.5	0.3
250	64 561	11.3	11.2	11.8	11.2	1.7	1.7	1.9	1.3
125	257 121	5.5	5.5	6.0	5.3	7.4	7.4	8.4	5.8

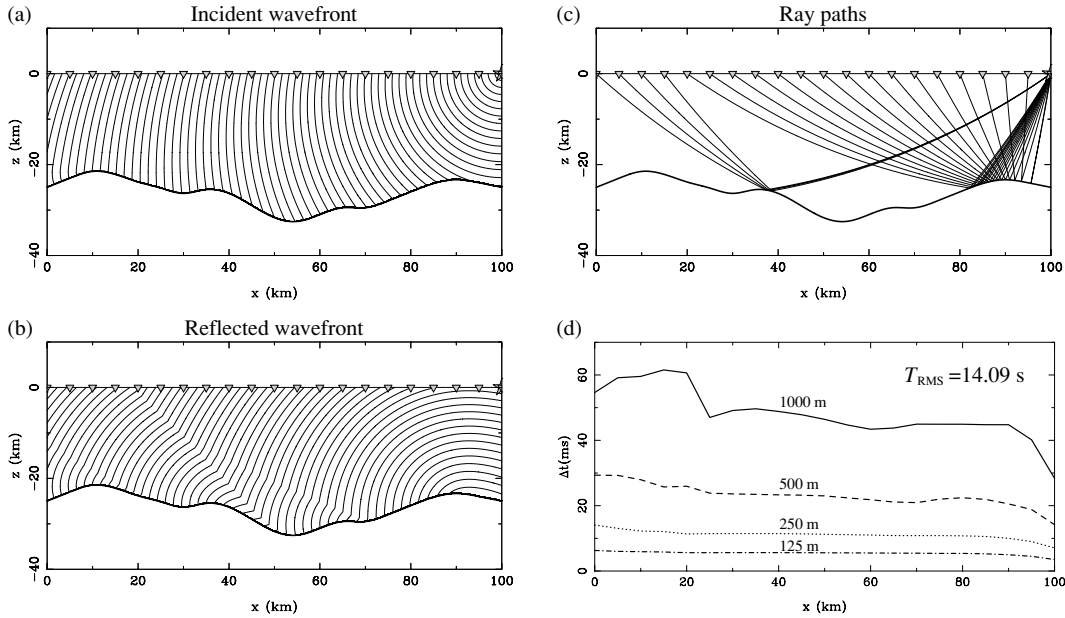


Figure 13. Same as Fig. 12 except with the source located at $(x, z) = (100, 0)$.

involves the location of the intersection point of the ray and interface, which uses an iterative Newton scheme. However, the accuracy of the intersection point can be controlled, so highly accurate solutions are possible. For the following comparisons we require that the ray-interface intersection point is accurate to 0.5 mm and the ray-receiver intersection point is accurate to 5 mm. Using such small values, we can be confident that the ray tracing traveltimes are, for all intents and purposes, exact.

The first example (Fig. 12) shows a wave front emanating from a point source and reflecting from an undulating interface. The ray paths associated with the receiver array reflect from two anticlinal structures that separate the upper and lower layers (Fig. 12c); the crossover of the first-arrival phase from one region of the reflector to the other is clear in the ray tracing diagram and is represented in the reflected wave front (Fig. 12b) by a propagating discontinuity. Comparison of FMM with the ray tracing solutions for four different grid sizes (Fig. 12d) shows that the FMM traveltimes converge to the exact solution as the grid size is reduced. With a 125 m grid spacing, the rms error (see Table 3) is only 5.5 ms, which is much smaller than would be required in most realistic applications. The CPU times for the FMM calculations are shown in Table 3; comparison of these times with those required for refracted wave fronts in continuous media (see Table 1) suggest that the reflection scheme is only approximately 40 per cent slower in this case. Using a source at $(x, z) = (100, 0)$ instead of $(x, z) = (0, 0)$ (see Fig. 13 and Table 3) results in almost identical rms errors and computation times.

If more than one interface is present, the local triangular mesh is still only used to suture the interface nodes of the reflector to adjacent

velocity nodes above the reflector. Any interface above the reflector will be approximated by a velocity gradient: interface nodes are only used to define the geometry of the reflector. Fig. 14 demonstrates the accuracy of FMM in such a situation. Here, the reflected wave front from the lower interface of a two interface model is tracked. Many of the ray paths (Fig. 14c) experience a significant deflection (particularly evident in the reflected branches) as they pass through the upper interface. The rms traveltimes errors associated with different grid sizes (see Table 3) are very similar to those of the single interface model and suggest that approximating a refracting interface with a sharp velocity gradient does not compromise the accuracy of the scheme. However, the spikes observed in the error plot (Fig. 14d) for the 1000 m grid (in particular) correspond to those rays that have experienced the greatest deflection at the upper interface, and it may be that in extreme cases the interface approximation for refractions results in significant error; this is examined more closely in the next example.

One common scenario that involves transmitted rays deviating significantly at layer boundaries occurs when rays turn back to the surface within a layer. Fig. 15 shows refracted rays turning back to the surface in a half-space below a single layer. The velocity contrast across the interface ($\sim 3.1 \text{ km s}^{-1}$) and the velocity gradient in the half-space ($\sim 0.1 \text{ s}^{-1}$) have been deliberately chosen to be large partly to help the ray tracer to locate turning rays and partly to test the method in extreme circumstances. Fig. 15(b) clearly shows that most paths are strongly deflected on both entry to and exit from the half-space. The traveltimes errors (Fig. 15c) fluctuate between adjacent stations but the general accuracy (Table 3) is similar to the

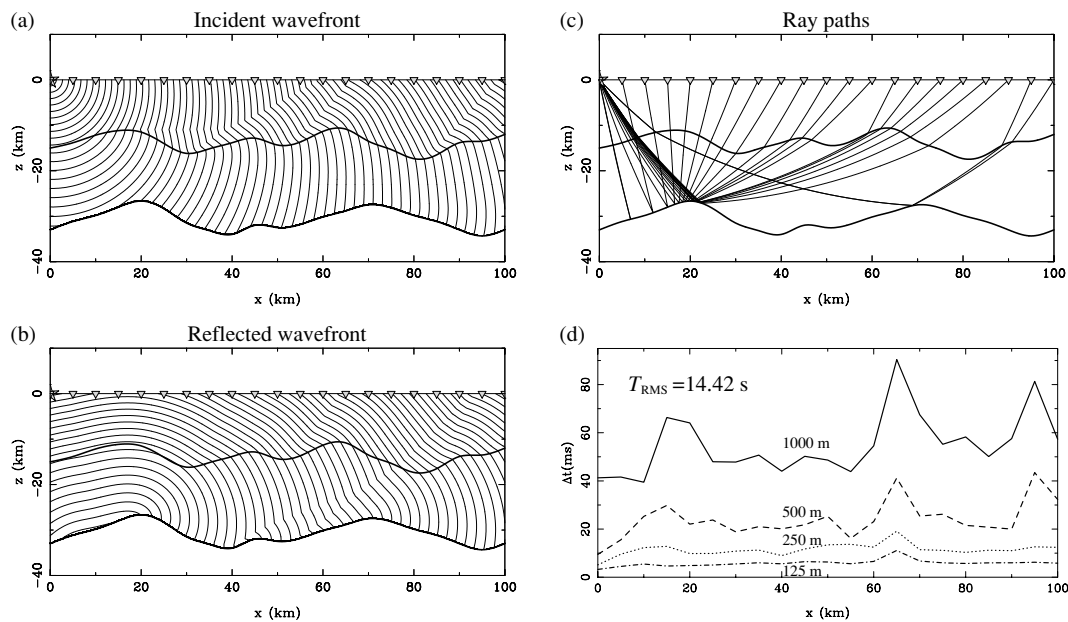


Figure 14. Comparison between the second-order FMM scheme and the shooting method of ray tracing in a medium containing two interfaces. The velocity at the surface is 4.5 km s^{-1} and the velocity gradients in the first and second layers are 0.04 and 0.05 s^{-1} , respectively. The velocity contrast across the top interface is approximately 1.2 km s^{-1} . Wave fronts are contoured at 0.3 s intervals.

reflected traveltimes of previous examples. Some of the observed fluctuation may be due to the use of second-order finite-difference operators in the presence of velocity discontinuities and the way in which the layer boundary is approximated.

We now investigate the effects of including lateral velocity variations within a layer in addition to variations in interface structure. At this point, the ray tracing method of Rawlinson *et al.* (2001a) can no longer be used for comparison because it requires velocity within a layer to vary with depth only. Therefore, we revert to the scheme used for analysing error in the Fig. 11 examples, which appeals to the stability of FMM and uses traveltimes calculated on a 50 m grid as a proxy for the exact traveltimes.

The example shown in Fig. 16(a) includes a strongly undulating interface within a medium that contains significant lateral velocity variation. The incident wave front is distorted by the velocity field before it impinges on the interface and the reflected wave front is dominated by the effects of the anticlinal interface structures. Each of these peaks act as a kind of virtual source for the reflections, as demonstrated by the almost spherical wave fronts that emanate from them. As these fronts propagate back towards the surface, they compete with the fronts generated by adjacent peaks for first-arrival status. Which wave front succeeds is easily seen in the ray path plots. An interesting feature of the rms traveltime error for the four different grid sizes (Table 4) is that it drops by more than half each time the grid spacing is halved. The final rms error estimate for the 125 m grid is only 3.0 ms .

The objective of the final example (Fig. 16b) is to emphasize the stability of the new scheme by applying it to a pathological velocity model. The model consists of three layers overlying a half-space; within each layer, velocity contrasts may be as large as 800 per cent, and interfaces have extreme curvature. In addition, the second layer below the surface pinches out at either end. The incident and reflected wave fronts are highly distorted by these structures, and many wave front gradient discontinuities may be observed. Note that the scheme is able to deal with a propagating shock in the incident field that impinges on the reflector (the bottom interface)

at approximately $x = 65 \text{ km}$. The rms traveltime error estimates (Table 4) for this example are significantly larger than the corresponding estimates for the previous example at larger grid spacing, but it is interesting to note that as the grid spacing is reduced, the errors converge. For example, the rms errors for the 1000 m grid are 62.6 and 104.3 ms for Figs 16(a) and (b), respectively, while for the 125 m grid they are only 3.0 and 3.5 ms , respectively.

4 DISCUSSION

First-arrival traveltimes are commonly used in seismic imaging, the appeal being that they are often straightforward to pick from a seismogram and simplify the modelling of the wavefield (i.e. only the path of minimum traveltime is required). In this regard, FMM is ideal because it is fast, unconditionally stable and guarantees that the first-arrival traveltime field is found. However, in some velocity media, the absolute first-arrival may be of such small amplitude that it will not be visible on a seismic record. An example of where this phenomenon might be observed is Fig. 11(b), where the first-arrival wave front is always strongly convex in the direction of ray propagation. As a result, the rays to all 21 receivers begin along practically identical paths and subsequent geometric spreading is extremely large. Even though this is an extreme example, similar behaviour can be observed in relatively simple media. For instance, first-arrival head waves, which can be generated in most layered media, carry very little energy.

Our extension of FMM to permit reflected phases to be calculated means that we are able to simulate a greater proportion of the total wavefield. In reflection and refraction seismology for example, reflected phases are often of far greater amplitude than first-arrival refractions, which make them relatively easy to pick (Rawlinson *et al.* 2001b). When non-planar interfaces are present, multiple two-point paths are usually possible; the FMM scheme will always locate the first-arrival reflection traveltime field in these cases. The examples we have illustrated in Section 3.2 show that first-arrival reflections tend to reflect from the elevated regions of the interfaces (the peaks).

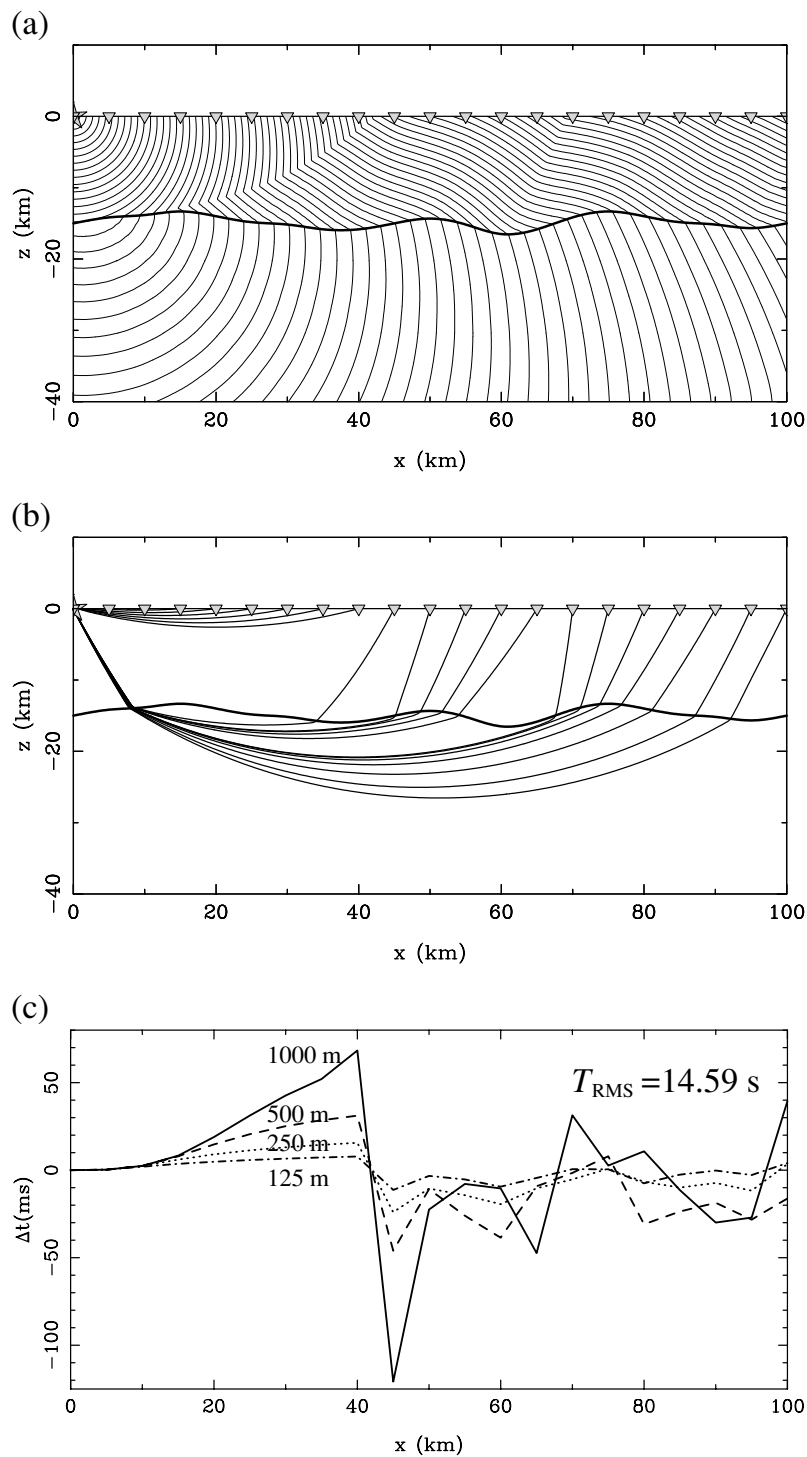


Figure 15. Comparison between the second-order FMM scheme and the shooting method of ray tracing for a refracted wave front. (a) First-arrival wave front (0.3 s contours); (b) first-arrival ray paths to all stations; (c) difference between FMM solution and shooting solution for four different grid sizes.

This means that the first-arrivals may only sample small portions of the interface and the wealth of later arriving reflection information is not considered.

Delprat-Jannaud & Lailly (1995) use a ray parameter based scheme to show that inclusion of later arrival reflection traveltimes in reflection tomography improves the imaging of complex structure. In particular, the inclusion of multi-valued traveltimes caused by triplications associated with synclines allows a much greater pro-

portion of the reflector to be illuminated. Clarke *et al.* (2001) extend the method to 3-D and demonstrate the superiority of using multi-arrival reflections compared with only first-arrival reflections in the imaging of a 3-D salt dome structure.

Clearly, our scheme would benefit from being able to track multipathing reflections and, more generally, triplications caused by continuous variations in wave speed. To date, a variety of methods have been proposed to solve the latter problem, including big ray

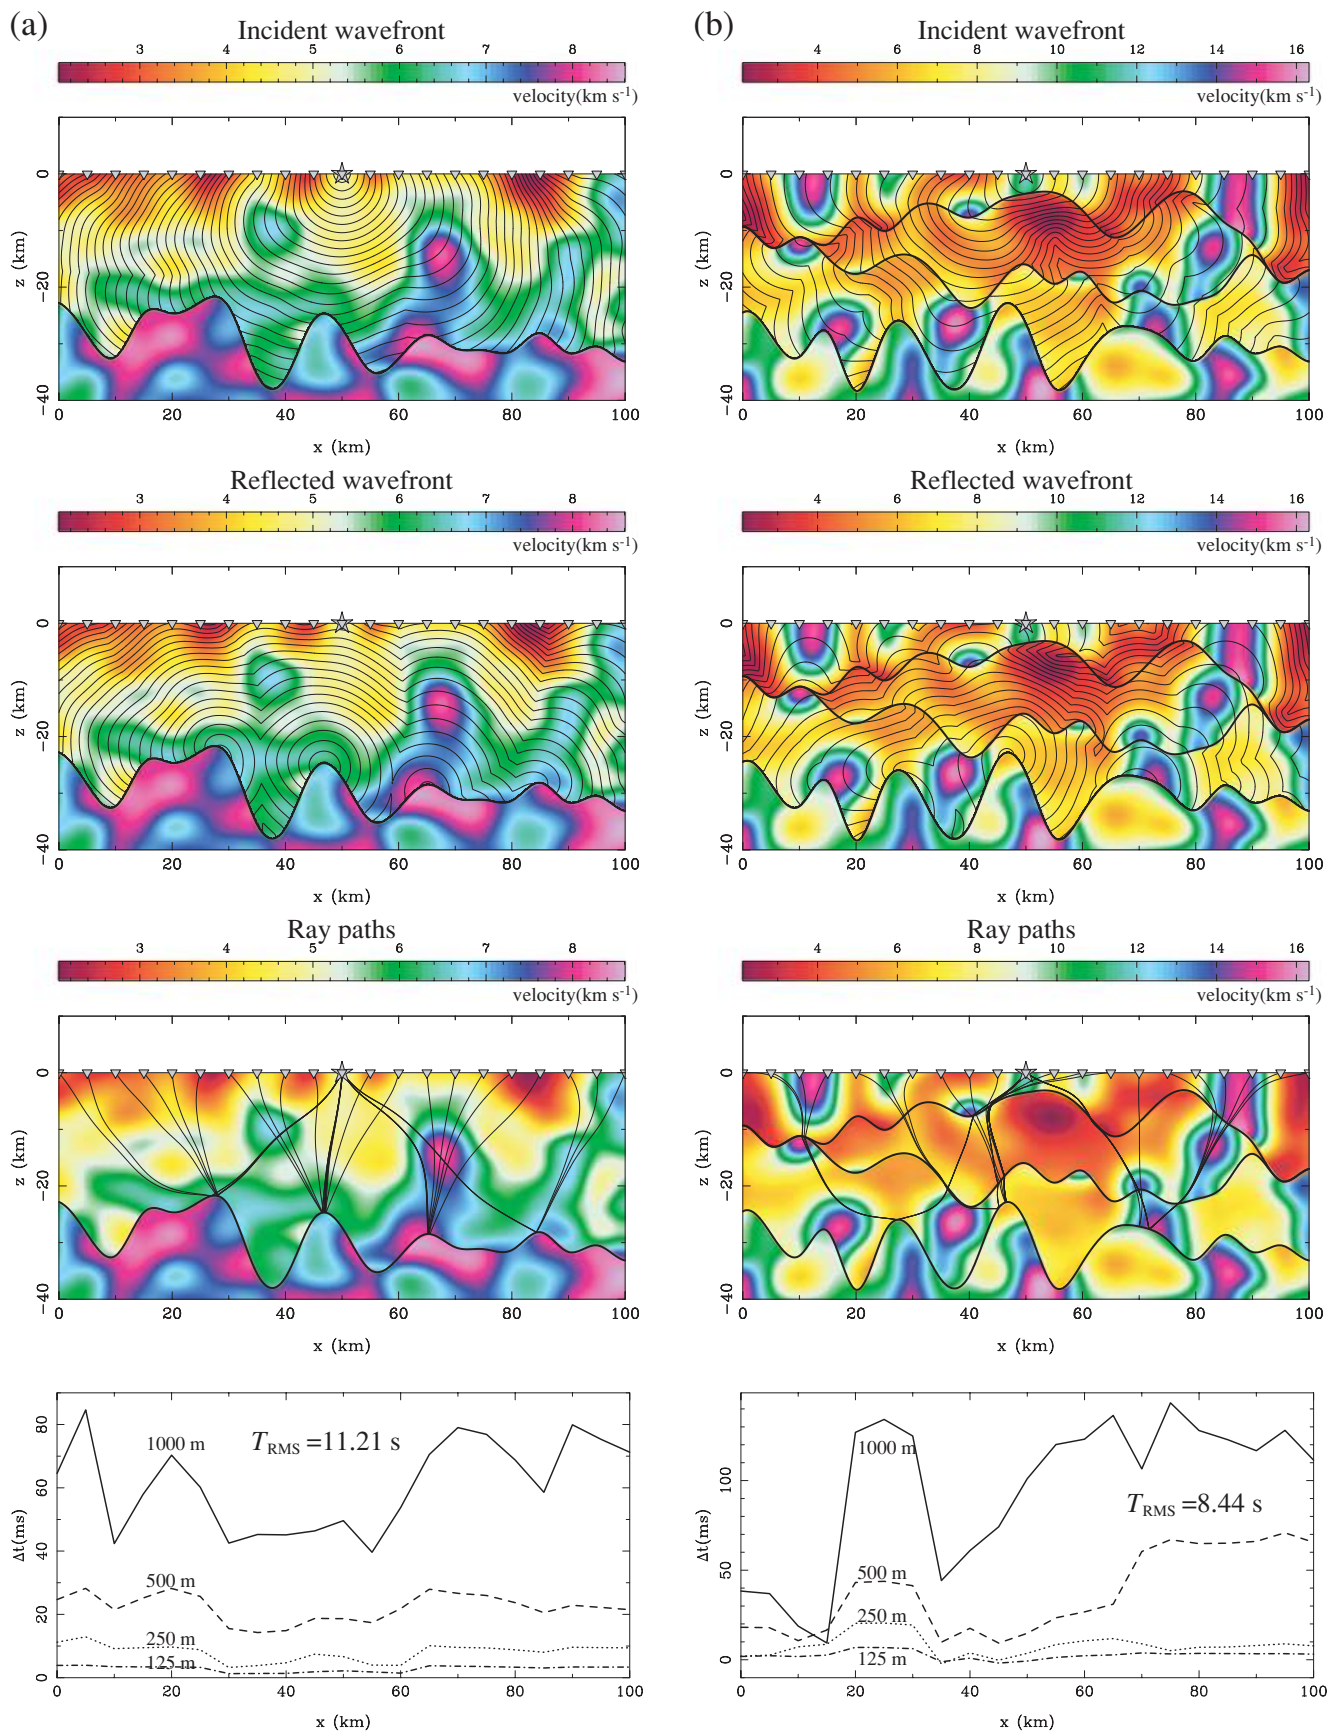


Figure 16. Reflection traveltimes in complex media. (a) Single interface model; (b) multiple interface model with velocity contrasts as large as 800 per cent. Wave fronts are contoured at 0.3 s intervals in both examples. Traveltimes error estimates are based on using FMM traveltimes from a 50 m grid.

Table 4. Summary of speed and accuracy tests for the two complex layered examples shown in Fig. 16.

Grid spacing (m)	Number of gridpoints	rms error (ms)		CPU time (s)	
		Fig. 16(a)	Fig. 16(b)	Fig. 16(a)	Fig. 16(b)
1000	4141	62.6	104.3	0.1	0.2
500	16 281	22.6	43.5	0.5	0.5
250	64 561	8.5	10.0	1.9	2.0
125	257 121	3.0	3.5	8.3	8.9

tracing (Abgrall 1999; Benamou 1996), dynamic surface evolution (Steinhoff *et al.* 2000) and wave front construction (Vinje *et al.* 1993; Lambaré *et al.* 1996). Recently, Fomel & Sethian (2002) introduced a scheme for computing phase space solutions of static Hamilton–Jacobi equations. Multiple arrivals in continuous media are found by solving partial differential equations on a grid using several ideas drawn from the fast marching method. The idea of exploiting phase space to find multi-valued solutions of partial differential equations that describe high frequency wave propagation has also been explored by Engquist *et al.* (2002) and Osher *et al.* (2002). Our scheme for layered media would benefit greatly from implementing these generalized techniques, but the computational cost would be significant.

Another way of tracking multiple arrivals in continuous media is to partition the multi-valued solution into a series of single valued solutions. A detailed examination of this approach is given by Benamou (1999); other interesting papers on the subject include those by Symes (1998) and Bevc (1997). The main advantage of this approach is that a first-arrival scheme is used to track wave fronts, so rapid computation can be achieved. To some extent, the scheme we propose is the layered equivalent of this approach, in that we partition the problem into separate computational domains and obtain later arriving phases (reflections) with a first-arrival solver.

Despite the problems associated with finding only the first-arrival component of phases as discussed above, the scheme we have proposed in this paper is practical for a variety of applications. These include reflection tomography, wide-angle tomography, local earthquake tomography, hypocentre relocation and migration. The results we have presented in Section 3 provide a general guideline for the accuracy of FMM in the presence of simple and complex structures. The required precision of model traveltimes is generally dependent on how accurately phases can be picked from seismograms. In crustal scale wide-angle tomography, for example, typical picking errors range from around 50 ms for first-arrivals to around 150 ms for later arriving reflections (e.g. Zelt & Smith 1992). Our results suggest that a 500 m computational grid with a second-order FMM would be sufficient for both phase types, although a 250 m grid may be required for the more accurately picked refraction arrivals.

Future development of the FMM scheme is possible on several fronts. As mentioned above, allowing phase-space computed multiple arrivals would increase the flexibility of the scheme but at significant computational cost. Partitioning the multi-valued solution as described by Benamou (1999) and others may be a better option as it is more compatible with the current scheme. One obvious area of development is to try and improve the accuracy of the scheme in the vicinity of the source. We demonstrate in Fig. 9 that if the large curvature of the wave front near the source is accounted for, the accuracy of the computed traveltimes can be improved by an order of magnitude in the constant velocity case. One way of doing this is to use a much finer computational grid in the vicinity of the source (e.g. Qian & Symes 2002). Alkhalifah & Fomel (2001) and Pica (1997) describe other schemes, based on spherical grids

and celerity, for overcoming errors generated in the vicinity of the source.

If local triangulation were permitted on the underside of an interface, then more phase types, such as later arriving refractions and multiples, could be calculated. The same principle for calculating reflections, i.e. propagate the wave front to all points on the interface and then reinitialize the scheme, could be extended to find multiple bounces or transmissions. Finally, the FMM scheme we have presented is only applicable to 2-D models; extension to 3-D would be a logical step. In continuous media, FMM is easily applied in 3-D, and the two step procedure for reflections is unchanged. The only major effort would be in extending the local irregular meshing scheme to 3-D, where tetrahedra rather than triangles are required.

FMM was developed in the field of computational mathematics (Sethian 1996) as a general method for tracking the evolution of monotonically advancing fronts. To date, its applications have been many and varied (Sethian 1999), with seismology forming only a small subset. Although FMM is considered to be a new scheme, it is interesting to note the similarities between it and pre-existing finite-difference grid-based methods developed in the seismic context. For example, van Trier & Symes (1991) use entropy-satisfying first-order upwind difference operators to numerically solve the eikonal equation, but also use an expanding rectangle as the computational front rather than the wave front. Qin *et al.* (1992) use the same difference operators as Vidale (1988) but calculates the traveltime field along the expanding wave front in a similar fashion to the narrow-band approach used by FMM. Podvin & Lecomte (1991) properly treat propagating wave front discontinuities by using Huygen's principle in the finite-difference approximation.

The FMM method we have presented introduces the concepts of local triangulation about an interface and using the set of interface nodes as the initial narrow band for evolving later arriving wave fronts. While this has not been done previously with FMM, Li & Ulrych (1993) use similar ideas to determine reflected and refracted traveltime fields in 2-D. In the vicinity of an interface, they use a local regridding technique that involves decomposing a cell containing an interface into several rectangular and triangular cells, so that the true interface shape is better represented on the computational grid. The incident wave front is tracked to the grid and then the scheme is reinitialized to track either the refracted or reflected wave front. Although this approach appears similar to the one presented here, entropy satisfying upwind operators are not used and Li & Ulrych (1993) only test the scheme on relatively simple models with planar interfaces.

We have noted that FMM encapsulates a number of ideas previously developed by the seismological community for grid-based solutions of the eikonal equation prior to the introduction of FMM. However, FMM is unique in that it uses entropy satisfying upwind difference operators along with a narrow-band approach to produce unconditionally stable solutions with minimum computational cost. The scheme we have presented preserves these properties and allows a larger proportion of the total seismic wavefield to be simulated.

5 CONCLUSIONS

We have presented an implementation of FMM in layered media, which allows first-arrival refraction and later arriving reflection traveltimes to be calculated. The scheme uses a local triangular mesh in the neighbourhood of the interface to suture irregularly spaced

interface nodes to adjacent velocity nodes and a two-stage evolution approach to propagate the incident and reflected wave fronts. In continuous media, we have shown that second-order upwind operators result in the best trade-off between accuracy and speed; higher-order operators will not necessarily produce more accurate solutions unless errors generated in the vicinity of the source are accounted for. In layered media, we compared the traveltimes produced by our FMM scheme with those produced by an accurate shooting method of ray tracing. Results suggest that the FMM scheme has both the speed and accuracy to be practical in applications that require the first-arrival or reflected wavefield to be simulated. Significantly, the extended scheme retains the desirable property of being unconditionally stable, as demonstrated by several examples in highly complex layered media. Future development of the scheme will be directed towards minimizing near source error, producing other classes of later arriving phases such as multiples, and extension to 3-D.

ACKNOWLEDGMENTS

The authors would like to thank Raul Madariaga and Gilles Lambaré for constructive reviews of an earlier version of this manuscript.

REFERENCES

- Abgrall, R., 1999. Big ray-tracing and eikonal solver on unstructured grids: application to the computation of a multivalued traveltime field in the Marmousi model, *Geophysics* **64**, 230–239.
- Afnimar & Koketsu, K., 2000. Finite difference traveltime calculation for head waves travelling along an irregular interface, *Geophys. J. Int.*, **143**, 729–734.
- Aki, K., Christofferson, A. & Husebye, E.S., 1977. Determination of the three-dimensional seismic structure of the lithosphere, *J. geophys. Res.*, **82**, 277–296.
- Alkhalifah, T. & Fomel, S., 2001. Implementing the fast marching eikonal solver: spherical versus Cartesian coordinates, *Geophys. Prospect.*, **49**, 165–178.
- Bartels, R.H., Beatty, J.C. & Barsky, B.A., 1987. *An introduction to Splines for Use in Computer Graphics and Geometric Modelling*, Morgan Kaufmann, Los Altos.
- Benamou, J.D., 1996. Big ray tracing: multivalued travel time field computation using viscosity solutions of the Eikonal equation, *J. Comp. Phys.*, **128**, 463–474.
- Benamou, J.D., 1999. Direct computation of multivalued phase space solutions for Hamilton–Jacobi equations, *Commun. Pure Appl. Math.*, **52**, 1443–1475.
- Bevc, D., 1997. Imaging complex structures with semirecursive Kirchhoff migration, *Geophysics* **62**, 577–588.
- Bijwaard, H. & Spakman, W., 2000. Non-linear global *P*-wave tomography by iterated linearized inversion, *Geophys. J. Int.*, **141**, 71–82.
- Bulant, P., 1996. Two-point ray tracing in 3-D, *Pure appl. Geophys.*, **148**, 421–447.
- Buske, S., 1999. 3-D prestack Kirchhoff migration of the ISO89–3D data set, *Pure appl. Geophys.*, **156**, 157–171.
- Cassell, B.R., 1982. A method for calculating synthetic seismograms in laterally varying media, *Geophys. J. R. astr. Soc.*, **69**, 339–354.
- Červený, V., 1987. Ray tracing algorithms in three-dimensional laterally varying layered structures, in *Seismic Tomography: With Applications in Global Seismology and Exploration Geophysics*, pp. 99–133, ed. Nolet, G., Reidel, Dordrecht.
- Červený, V., 2001. *Seismic Ray Theory*, Cambridge University Press, Cambridge.
- Cheng, N. & House, L., 1996. Minimum traveltime calculations in 3-D graph theory, *Geophysics* **61**, 1895–1898.
- Chopp, D.L., 2001. Some improvements of the fast marching method, *SIAM J. Sci. Comput.*, **23**, 230–244.
- Clarke, R.A., Alazard, B., Pelle, L., Sinoquet, D., Lailly, P., Delprat-Jannaud, F. & Jannaud, L., 2001. 3D traveltime reflection tomography with multivalued arrivals, *71st Ann. Int. Mtg. Soc. Expl. Geophys., Expanded Abstracts*, pp. 1875–1878.
- Darbyshire, F.A., Bjarnason, I.J., White, R.S. & Flórenz, O.G., 1998. Crustal structure above the Iceland mantle plume, imaged by the ICEMELT refraction profile, *Geophys. J. Int.*, **135**, 1131–1149.
- Delprat-Jannaud, F. & Lailly, P., 1995. Reflection tomography: how to handle multiple arrivals?, *J. geophys. Res.*, **100**, 703–715.
- Dziewonski, A.M. & Woodhouse, J.H., 1987. Global images of the Earth's interior, *Science*, **236**, 37–48.
- Eberhart-Phillips, D. & Michael, A.J., 1993. Three-dimensional velocity structure, seismicity, and fault structure in the Parkfield Region, central California, *J. geophys. Res.*, **98**, 15 737–15 758.
- Engquist, B., Runborg, O. & Tornberg, A.-K., 2002. High-frequency wave propagation by the segment projection method, *J. Comp. Phys.*, **178**, 373–390.
- Epili, D. & McMechan, G.A., 1996. Implementation of 3-D prestack Kirchhoff migration, with application to data from the Ouachita frontal thrust zone, *Geophysics* **61**, 1400–1411.
- Ettrich, N. & Gajewski, D., 1996. Wave front construction in smooth media for prestack depth migration, *Pageophys.*, **148**, 481–502.
- Fischer, R. & Lees, J.M., 1993. Shortest path ray tracing with sparse graphs, *Geophysics* **58**, 987–996.
- Fomel, S. & Sethian, J.A., 2002. Fast-phase space computation of multiple arrivals, *Proc. Natl Acad. Sci., USA*, **99**, 7329–7334.
- Gorbatov, A., Widiyantoro, S., Fukao, Y. & Gordeev, E., 2000. Signature of remnant slabs in the North Pacific from *P*-wave tomography, *Geophys. J. Int.*, **142**, 27–36.
- Gorbatov, A., Fukao, Y., Widiyantoro, S. & Gordeev, E., 2001. Seismic evidence for a mantle plume oceanwards of the Kamchatka–Aleutian trench junction, *Geophys. J. Int.*, **146**, 282–288.
- Graeber, F.M. & 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, *J. geophys. Res.*, **104**, 20 237–20 256.
- Graeber, F.M., Houseman, G.A. & Greenhalgh, S.A., 2002. Regional teleseismic tomography of the western Lachlan Orogen and the Newer Volcanic Province, southeast Australia, *Geophys. J. Int.*, **149**, 249–266.
- Gray, S.H. & May, W.P., 1994. Kirchhoff migration using eikonal equation traveltimes, *Geophysics*, **59**, 810–817.
- Grechka, V.Y. & McMechan, G.A., 1996. 3-D two-point ray tracing for heterogeneous, weakly transversely isotropic media, *Geophysics* **61**, 1883–1894.
- Gudmundsson, O., Brandsdóttir, B., Menke, W. & Sigvaldason, G.E., 1994. The crustal magma chamber of the Katla volcano in south Iceland revealed by 2-D seismic undershooting, *Geophys. J. Int.*, **119**, 277–296.
- Hole, J.A. & Zelt, B.C., 1995. 3-D finite-difference reflection travel times, *Geophys. J. Int.*, **121**, 427–434.
- Hole, J.A., Clowes, R.M. & Ellis, R.M., 1992. Interface inversion using broadside seismic refraction data and three-dimensional traveltime calculations, *J. geophys. Res.*, **97**, 3417–3429.
- Inoue, H., Fukao, Y., Tanabe, K. & Ogata, Y., 1990. Whole mantle *P*-wave travel time tomography, *Phys. Earth planet. Inter.*, **59**, 294–328.
- Julian, B.R. & Gubbins, D., 1977. Three-dimensional seismic ray tracing, *J. Geophys.*, **43**, 95–113.
- Lambaré, G., Lucio, P.S. & Hanyga, A., 1996. Two-dimensional multivalued traveltime and amplitude maps by uniform sampling of a ray field, *Geophys. J. Int.*, **125**, 584–598.
- Leidenfrost, A., Ettrich, N., Gajewski, D. & Kosloff, D., 1999. Comparison of six different methods for calculating traveltimes, *Geophys. Prospect.*, **47**, 269–297.
- Li, X.G. & Ulrych, T.J., 1993. Traveltime computation in discrete heterogeneous layered media, *Jrnl. Seismic Explor.*, **2**, 305–318.
- Lucio, P.S., Lambaré, G. & Hanyga, A., 1996. 3D multivalued travel time and amplitude maps, *Pageoph*, **148**, 449–479.

- Morelli, A., 1993. Teleseismic tomography: core–mantle boundary, in *Seismic Tomography: Theory and Practice*, pp. 163–189, eds Iyer, H.M. & Hirahara, K., Chapman & Hall, London.
- Moser, T.J., 1991. Shortest path calculation of seismic rays, *Geophysics* **56**, 59–67.
- Nakanishi, I. & Yamaguchi, K., 1986. A numerical experiment on nonlinear image reconstruction from first-arrival times for two-dimensional island arc structure, *J. Phys. Earth*, **34**, 195–201.
- Oncescu, M.C., Burlacu, V., Anghel, M. & Smalbergher, V., 1984. Three-dimensional *P*-wave velocity image under the Carpathian Arc, *Tectonophysics*, **106**, 305–319.
- Osher, S., Cheng, L.-T., Kang, M., Shim, H. & Tsai, Y.-H., 2002. Geometric optics in a phase-space-based level set and Eulerian framework, *J. Comp. Phys.*, **179**, 622–648.
- Pereyra, V., Lee, W.H.K. & Keller, H.B., 1980. Solving two-point seismic-ray tracing problems in a heterogeneous medium, *Bull. seism. Soc. Am.*, **70**, 79–99.
- Pica, A., 1997. Fast and accurate finite-difference solutions of the 3D eikonal equation parametrized in celerity, *67th Ann. Int. Mtg. Soc. Expl. Geophys., Expanded Abstracts*, pp. 1774–1777.
- Podvin, P. & Lecomte, I., 1991. Finite difference computation of traveltimes in very contrasted velocity models: a massively parallel approach and its associated tools, *Geophys. J. Int.*, **105**, 271–284.
- Popovici, A.M. & Sethian, J.A., 2002. 3-D imaging using higher order fast marching traveltimes, *Geophysics* **67**, 604–609.
- Prothero, W.A., Taylor, W.J. & Eickemeyer, J.A., 1988. A fast, two-point, three-dimensional raytracing algorithm using a simple step search method, *Bull. seism. Soc. Am.*, **78**, 1190–1198.
- Qian, J. & 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. & Schuster, G.T., 1992. Finite-difference solution of the eikonal equation along expanding wavefronts, *Geophysics* **57**, 478–487.
- Rawlinson, N., Houseman, G.A. & Collins, C.D.N., 2001a. Inversion of seismic refraction and wide-angle reflection traveltimes for 3-D layered crustal structure, *Geophys. J. Int.*, **145**, 381–401.
- Rawlinson, N., Houseman, G.A., Collins, C.D.N. & Drummond, B.J., 2001b. New evidence of Tasmania's tectonic history from a novel seismic experiment, *Geophys. Res. Lett.*, **28**, 3337–3340.
- Riahi, M.A. & Juhlin, C., 1994. 3-D interpretation of reflected arrival times by finite-difference techniques, *Geophysics* **59**, 844–849.
- Sambridge, M.S. & Kennett, B.L.N., 1990. Boundary value ray tracing in a heterogeneous medium: a simple and versatile algorithm, *Geophys. J. Int.*, **101**, 157–168.
- Sena, A.G. & Toksöz, M.N., 1993. Kirchhoff migration and velocity analysis for converted and nonconverted waves in anisotropic media, *Geophysics* **58**, 265–276.
- Sethian, J.A., 1996. A fast marching level set method for monotonically advancing fronts, *Proc. Nat. Acad. Sci.*, **93**, 1591–1595.
- Sethian, J.A., 1999. *Level Set Methods and Fast Marching Methods*, Cambridge University Press, Cambridge.
- Sethian, J.A., 2001. Evolution, implementation, and application of level set and fast marching methods for advancing fronts, *J. Comp. Phys.*, **169**, 503–555.
- Sethian, J.A. & Popovici, A.M., 1999. 3-D traveltimes computation using the fast marching method, *Geophysics* **64**, 516–523.
- Spakman, W., 1991. Delay-time tomography of the upper mantle below Europe, the Mediterranean and Asia Minor, *Geophys. J. Int.*, **107**, 309–332.
- Steinoff, J., Fan, M. & Wang, L., 2000. A new Eulerian method for the computation of propagating short acoustic and electromagnetic pulses, *J. Comp. Phys.*, **157**, 683–706.
- Symes, W.W., 1998. A slowness matching finite difference method for traveltimes beyond transmission caustics, *68th Ann. Int. Mtg. Soc. Expl. Geophys., Expanded Abstracts*, pp. 1945–1948.
- Thurber, C.H., 1983. Earthquake locations and three-dimensional crustal structure in the Coyote Lake area, central California, *J. geophys. Res.*, **88**, 8226–8236.
- Um, J. & Thurber, C., 1987. A fast algorithm for two-point seismic ray tracing, *Bull. seism. Soc. Am.*, **77**, 972–986.
- van der Hilst, R.D., Widiyantoro, S. & Engdahl, E.R., 1997. Evidence for deep mantle circulation from global tomography, *Nature*, **386**, 578–584.
- van Trier, J. & Symes, W.W., 1991. Upwind finite-difference calculation of traveltimes, *Geophysics* **56**, 812–821.
- Vidale, J.E., 1988. Finite-difference calculations of traveltimes, *Bull. seism. Soc. Am.*, **78**, 2062–2076.
- Vidale, J.E., 1990. Finite-difference calculations of traveltimes in three dimensions, *Geophysics* **55**, 521–526.
- Vinje, V., Iversen, E. & Gjøystdal, H., 1993. Traveltime and amplitude estimation using wavefront construction, *Geophysics* **58**, 1157–1166.
- Vinje, V., Åstebøl, K., Iversen, E. & Gjøystdal, H., 1999. 3-D ray modelling by wavefront construction in open models, *Geophys. Prospect.*, **64**, 1912–1919.
- Widiyantoro, S. & van der Hilst, R., 1997. Mantle structure beneath Indonesia inferred from high-resolution tomographic imaging, *Geophys. J. Int.*, **130**, 167–182.
- Zelt, B.C., Ellis, R.M., Zelt, C.A., Hyndman, R.D., Lowe, C., Spence, G.D. & Fisher, M.A., 2001. Three-dimensional crustal velocity structure beneath the Strait of Georgia, British Columbia, *Geophys. J. Int.*, **144**, 695–712.
- Zelt, C.A. & Smith, R.B., 1992. Seismic traveltimes inversion for 2-D crustal velocity structure, *Geophys. J. Int.*, **108**, 16–34.
- Zhao, D., Hasegawa, A. & Horiuchi, S., 1992. Tomographic imaging of *P* and *S* wave velocity structure beneath Northeastern Japan, *J. geophys. Res.*, **97**, 19 909–19 928.