

Multiple reflection and transmission phases in complex layered media using a multistage fast marching method

N. Rawlinson¹ and M. Sambridge¹

ABSTRACT

Traditional grid-based eikonal schemes for computing traveltimes are usually confined to obtaining first arrivals only. However, later arrivals can be numerous and of greater amplitude, making them a potentially valuable resource for practical applications such as seismic imaging. The aim of this paper is to introduce a grid-based method for tracking multivalued wavefronts composed of any number of reflection and refraction branches in layered media. A finite-difference eikonal solver known as the fast marching method (FMM) is used to propagate wavefronts from one interface to the next. By treating each layer that the wavefront enters as a separate computational domain, one obtains a refracted branch by reinitializing FMM in the adjacent layer and a reflected branch by reinitializing FMM in the incident layer.

To improve accuracy, a local grid refinement scheme is used in the vicinity of the source where wavefront curvature is high. Several examples are presented which demonstrate the viability of the new method in highly complex layered media. Even in the presence of velocity variations as large as 8:1 and interfaces of high curvature, wavefronts composed of many reflection and transmission events are tracked rapidly and accurately. This is because the scheme retains the two desirable properties of a single-stage FMM: computational speed and stability. Local grid refinement about the source also can increase accuracy by an order of magnitude with little increase in computational cost.

INTRODUCTION

The complex nature of many seismic wavetrains can usually be attributed to elastic wave energy multipathing between source and receiver, caused by both continuous and discontinuous variations in seismic wavespeed. This paper focuses

on the latter case, when reflections and transmissions at interfaces generate a variety of phases. A vast array of techniques are currently available for modeling the propagation of (high-frequency) seismic waves in laterally heterogeneous media, but most are only capable of consistently locating first arrivals and/or single reflection phases between source and receiver.

Ray tracing has traditionally been the method of choice for tracking the path taken by seismic energy between two points. Shooting methods (e.g., Julian and Gubbins, 1977; Červený, 1987; Sambridge, 1990; Bulant, 1996; Rawlinson et al., 2001) formulate the ray equation as an initial value problem and iteratively adjust the ray take-off angle until a source–receiver path is located. Bending methods (Julian and Gubbins, 1977; Um and Thurber, 1987; Grechka and McMechan, 1996) iteratively adjust the geometry of an initial arbitrary path that joins the source and receiver until it satisfies Fermat's principle. Many ray-tracing methods have been developed to find first arrivals in continuous media (e.g., Julian and Gubbins, 1977; Pereyra et al., 1980; Prothero et al., 1988) and to locate refraction and reflection arrivals in layered media (e.g., Farra and Madariaga, 1988; Lutter and Nowack, 1990; Zelt and Smith, 1992; Zhao et al., 1997; Rawlinson et al., 2001). Although ray tracing can be rapid and highly accurate, it is not robust and may fail to locate a two-point path even in the presence of relatively simple velocity structures, e.g., near critical refractions in horizontal stratification (Kim and Baag, 2002).

An alternative approach to ray tracing is to seek finite-difference solutions of the eikonal equation throughout a gridded velocity field (e.g., Vidale, 1988, 1990; Podvin and Lecomte, 1991; van Trier and Symes, 1991; Hole and Zelt, 1995; Kim and Cook, 1999; Afnimar and Koketsu, 2000; Kim, 2002; Qian and Symes, 2002). While this approach can be fast, accurate, and robust, it only calculates the absolute first-arrival traveltimes to all points in a continuous medium.

Relatively few studies have used eikonal solvers to compute reflection and refraction traveltimes in layered media. To find reflections, Podvin and Lecomte (1991) and Riahi and Juhlin (1994) track the first-arrival traveltimes from both the source and the receiver to the entire interface. Fermat's

Manuscript received by the Editor March 31, 2003; revised manuscript received November 29, 2003.

¹Australian National University, Research School of Earth Sciences, Centre for Advanced Data Inference, Canberra, ACT 0200, Australia. E-mail: nick@rses.anu.edu.au; malcolm@rses.anu.edu.au.

© 2004 Society of Exploration Geophysicists. All rights reserved.

principle of stationary time is then used to locate reflection points along the interface. The advantage of this approach is that more than one reflection path between two points can be identified, but the clear drawback is that a traveltime field needs to be computed for each source and receiver. Hole and Zelt (1995) overcome this problem by assuming that the incident wavefront from the source and the interface are sufficiently smooth to validate a local planar approximation. Thus, the reflected traveltimes to the nodes that are adjacent to the interface can be determined using only the depth to the reflector, the normal vector to the reflector, and the direction vectors of the incident and reflected rays (from Snell's law). Strong wavefront or reflector curvature will degrade the accuracy of this scheme.

Li and Ulrych (1993) compute reflected and refracted traveltimes in two dimensions by using a local regridding technique to decompose 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 traveltime field is computed using the scheme of Vidale (1988), and the reflected traveltime field is obtained by reinitializing the computational front from the point of minimum traveltime on the interface. Although this scheme has some similarities to the one we present, it is tested only on relatively simple models with planar interfaces.

Shortest path ray tracing is another popular method for determining first arrivals to all points of a gridded velocity field (e.g., Nakanishi and Yamaguchi, 1986; Moser, 1991; Fischer and Lees, 1993; Cheng and House, 1996). Rather than solve a differential equation, a network or graph is formed by connecting neighboring nodes with traveltime path segments. The problem is then to find the path of minimum traveltime between source and receiver through the network, which may be solved using Dijkstra-like algorithms. By definition, shortest path ray tracing finds first-arrival traveltimes, but it is possible to impose constraints on the path so that some other arrivals such as reflections can be determined. Moser (1991) demonstrates a method for finding reflections which requires the shortest path to visit a specified set of nodes that lie on the interface. According to Leidenfrost et al. (1999), finite-difference methods offer a better compromise between speed and accuracy than does the shortest path ray-tracing method, but the latter is often considered more numerically stable (see Cheng and House, 1996).

In this paper, we present an elegant and robust grid-based approach for tracking wavefronts composed of any number of reflection and refraction branches in strongly heterogeneous layered media. At its core, the new scheme uses the fast marching method, or FMM (Sethian, 1996), which tracks the evolution of wavefronts via finite-difference solution of the eikonal equation. As discussed previously, this class of scheme is restricted to finding only the absolute first arrival in continuous media, but FMM distinguishes itself by combining both speed and unconditional stability.

To track arrivals composed of multiple reflection and refraction branches, 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. At this stage, FMM is halted and we are left with a narrow band of traveltime values defined along the interface. From here, a refracted branch can be tracked by reinitializing FMM from the narrow band into the adjacent layer, and a reflected branch can be obtained by reinitializing FMM in the

incident layer. Wavefronts composed of any number of reflection and refraction events can therefore be built up by using this multistage approach. A triangulation routine is used to locally suture interface nodes to neighboring velocity nodes in order to facilitate the tracking of wavefronts to and from the interface. To improve accuracy, local grid refinement about the source is implemented.

After describing the multistage FMM, we present several examples that explore the stability and accuracy of the method in simple and complex 2D layered media. We aim to demonstrate that the scheme retains the desirable characteristics of the single-stage FMM and is a practical way of tracking multiple phases in layered media using a grid-based method.

METHOD

The fast marching method

FMM is a grid-based numerical scheme for tracking the evolution of monotonically advancing interfaces by seeking finite-difference solutions to the eikonal equation $|\nabla_{\mathbf{x}}T| = s(\mathbf{x})$, where T is traveltime and $s(\mathbf{x})$ is slowness as a function of position \mathbf{x} . To date, it has been applied to a wide variety of problems in the physical sciences, including seismic wave propagation, photolithographical development, geodesics, deposition of sediments, medical imaging, and optimal path planning (see Sethian, 1999; Sethian and Popovici, 1999; Popovici and Sethian, 2002). Following is a brief description of the method; for more details, refer to Sethian (1996, 1999).

To achieve unconditional stability and computational speed, FMM combines solution of the eikonal equation using entropy satisfying upwind finite-difference operators with a narrow-band approach for the ordered update of gridpoints. A significant obstacle faced by all eikonal solvers is that the first-arrival wavefront may be discontinuous in gradient. When this occurs, it is difficult to solve the eikonal equation because it requires that $\nabla_{\mathbf{x}}T$ be defined. One way to overcome this problem is to solve the viscous version of the eikonal equation, which smooths out discontinuities; the limit of smooth solutions is a weak solution that corresponds to the first-arriving wavefront. It turns out that the viscous limit solution can also be obtained by solving the nonviscous eikonal equation using entropy-satisfying operators. The entropy condition states that the first-arrival wavefront can only pass through a point once and must be strictly enforced to guarantee unconditional stability.

Our entropy-satisfying upwind scheme has been used by a number of authors (e.g., Sethian and Popovici, 1999; Chopp, 2001; Popovici and Sethian, 2002). For 2D media it may be written as

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

where (i, j) are grid increment variables in (x, z) , and the integer variables a, b, c, d define the order of accuracy of the upwind finite-difference operator used in each of the four cases. In the case of $D^{-x}T_i$, the first- and second-order accurate operators are

$$D_1^{-x}T_i = \frac{T_i - T_{i-1}}{\delta x},$$

$$D_2^{-x}T_i = \frac{3T_i - 4T_{i-1} + T_{i-2}}{2\delta x}, \quad (2)$$

where δx is the grid spacing in x . In the examples that follow, we compare the accuracy of a purely first-order scheme with a mixed-order scheme that preferentially uses D_2 operators but reverts to D_1 operators when the required upwind traveltimes are unavailable. For convenience, we refer to this mixed-order scheme as a second-order scheme. Strictly speaking, the unconditional stability of FMM has only been proven for the first-order scheme; however, our results show that the second-order method is effective even in highly complex media.

The narrowband concept used by FMM to determine the order in which gridpoints are updated is illustrated in Figure 1. All gridpoints are labeled as 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 1 with Alive points only; Far points 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 neighboring Far points as Close. Finally, all Close points adjacent to the new Alive point have their traveltimes updated using equation 1. 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 $O(M \log M)$, where M is the total number of gridpoints.

Multistage FMM for layered media

The basic FMM method outlined above is designed for a regularly gridded continuous velocity field. In a layered medium with undulating interfaces, a purely regular grid is not suitable. Instead, we define velocity on a regular grid and then use an adaptive triangular mesh to suture interface nodes to neighboring velocity nodes, as illustrated in Figure 2. 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 minimized.




Figure 1. Principle of the narrowband method. Close points have trial values computed from Alive points, which lie upwind and have correct traveltime values. Far points lie downwind and are yet to have the wavefront pass through them.

Equation 1 is used to update gridpoints in the regular mesh, and a first-order accurate scheme for triangular elements is used to evolve the wavefront through the irregular mesh. To devise such a scheme, consider the triangular element shown in Figure 3a and assume that traveltimes to points A (T_A) and B (T_B) are known. 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 Sethian (1999) for a derivation]

$$(a^2 + b^2 - 2\mathbf{a} \cdot \mathbf{b})t^2 + 2u(\mathbf{a} \cdot \mathbf{b} - b^2)t + (b^2u^2 - s_O^2[a^2b^2 - (\mathbf{a} \cdot \mathbf{b})^2]) = 0, \quad (3)$$

where s_O is slowness at O, $u = T_B - T_A$, \mathbf{a} and \mathbf{b} are displacement vectors (see Figure 3a), $a = |\mathbf{a}|$, and $b = |\mathbf{b}|$. Note that equation 3 avoids the trigonometric functions used in the equivalent expression given by Sethian (1999).

Provided that the traveltime vector at point O lies between points A and B and $u < t$, equation 3 can be legitimately used to produce a first-order accurate solution $T_O = t + T_A$. If these conditions are not satisfied, then the solution is given by $T_O = \min\{bs_O + T_A, as_O + T_B\}$, which is only first-order accurate if the traveltime gradient is parallel to either vector \mathbf{a} or \mathbf{b} . If point O is connected to more than one triangle, then the traveltime prediction for each of these triangles needs to be calculated and the minimum chosen as the correct solution.

To guarantee that points within the triangular mesh are updated with first-order accuracy, adjacent mesh lines emanating from a node must be acute. To understand why this restriction exists, consider Figure 3b, which shows a wavefront impinging on a triangle with an obtuse angle at point O. If the wavefront normal lies within the zone subtended by θ , then clearly the wavefront will pass through both A and B before reaching point O; however, if the wavefront normal lies within the zones defined by ϕ_1 or ϕ_2 , then the wavefront will only pass through A or B, respectively, before reaching point O. In this case, first-order accuracy occurs only if the wavefront normal is parallel with AO or BO.




Figure 2. An adaptive triangular mesh is used to locally suture the irregular interface nodes to the regular nodes of the velocity grid (circles). The interface is plotted as a piecewise linear curve between the interface nodes. For visual purposes, the grid is much coarser than would be used in practical applications. A single source and 21 receivers located on the surface are denoted by a star and triangles, respectively.

For the local adaptive triangulation scheme illustrated in Figure 2, obtuse triangles can be generated, but only when the interface cuts a corner off a cell. An example of this effect is shown in Figure 3c, where the angle at point O is obtuse as a result of the low gradient of the interface joining OC. To fix this problem, we can use a different splitting strategy, which involves replacing mesh line AC with BO (see Figure 3d). Thus, to update point C, the mesh shown in Figure 3c is used; to up-

date point O, the mesh shown in Figure 3d is used. A more detailed description of the meshing routine and the first-order FMM for triangular elements 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 continue propagating through to the adjacent layer, the




Figure 3. Updating traveltimes in a triangulated domain. (a) If the traveltime gradient is assumed to be constant within a triangle, point O can be updated using known traveltimes at points A and B. (b) The presence of obtuse triangles can result in traveltimes that are not first-order accurate. (c). (d) Dealing with obtuse triangles by using a different splitting strategy.




Figure 4. Schematic of how a reflected or refracted wavefront can be tracked by partitioning layers into separate computational domains in which FMM is applied. (a) Incident wavefront generated from a point source; (b) narrow band defined by the set of interface nodes; (c) reflected wavefront tracked; (d) refracted wavefront tracked. In both (c) and (d), FMM reinitializes from the interface node with minimum traveltime.

wavefront is only tracked as far as the interface, which is treated as one of the 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 terminates. The next step is to track a transmitted or reflected wavefront using the traveltimes information we have obtained as a starting point. Interestingly, this can be done by using interface node traveltimes only, because they store sufficient information for first-order accurate estimates of traveltimes to neighboring nodes. The complete set of interface nodes is therefore used as the starting narrow band for the next FMM stage (Figure 4b). The fact that this narrow band will not in general conform to the shape of the wavefront does not matter, provided it is understood that the wavefront cannot reflect from or transmit through the interface more than once.

From the initial narrow band described by the set of interface nodes, a reflected wavefront can be tracked by setting all other nodes within the incident layer to Far and restarting FMM (see Figure 4c). A refracted wavefront is tracked in the same fashion, except that only the nodes in the adjacent layer are set to Far so the wavefront continues to propagate into the next layer (see Figure 4d) rather than reflect back into the incident layer. This 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 a wavefront can experience is limited by CPU time only. In practice, the total energy carried by the wavefront decreases monotonically with time, so one may not need to track phases composed of large numbers of refraction and reflection events to model real observations.

As previously noted, when a wavefront is required to transmit through an interface, it can be stopped at the interface and then reinitialized using the interface nodes as the initial narrow band. Our method allows for two further possibilities: (1) FMM is not reinitialized at the interface and the wavefront propagates directly through to the adjacent layer (i.e., both layers lie within the same computational domain); (2) the interface and associated irregular grid are excluded and the layer boundary is instead represented by a sharp velocity gradient. The first option means that the wavefront is permitted to pass through an interface more than once during a single FMM step. The second option allows layer pinch-outs to be represented without superimposing irregular meshes from adjacent interfaces.

Error minimization in the source neighborhood

A point source is an upwind singularity of the traveltimes field and can be a major contributor to the overall error of FMM because of high wavefront curvature and first-order accuracy

Table 1. Summary of speed and accuracy tests for the reflection example shown in Figure 6. CPU times are for a SunBlade 150 and include the time taken to read in the velocity field and interface surfaces.

Grid spacing (m)	Number of gridpoints	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	253.0	184.9	50.6	10.3	0.1	0.1	0.1	0.1
500	16 281	150.7	116.8	23.5	2.8	0.3	0.3	0.3	0.4
250	64 561	86.9	69.9	11.3	0.8	1.1	1.2	1.3	1.4
125	257 121	49.0	40.5	5.5	0.3	4.8	4.9	5.4	5.5

O = Order. R = Refined grid about source.




Figure 5. Implementation of source-grid refinement. When the narrow band (thick grey line) reaches the boundary of the refined grid, it is mapped onto the coarse grid (triangles) before continuing to evolve. Refined gridpoints that do not coincide with the coarse gridpoints are denoted by circles. The source is denoted by a large grey circle.

in its vicinity. This problem has been recognized and addressed in a variety of ways. For example, Alkhalifah and Fomel (2001) use spherical grids centered on the source point; Kim and Cook (1999) use local grid refinement in the source neighborhood;

Qian and Symes (2002) use adaptive gridding with grid refinement and coarsening based on a posteriori error estimation.

Our approach for minimizing near-source error is based on specifying a finely spaced grid in the vicinity of the source and

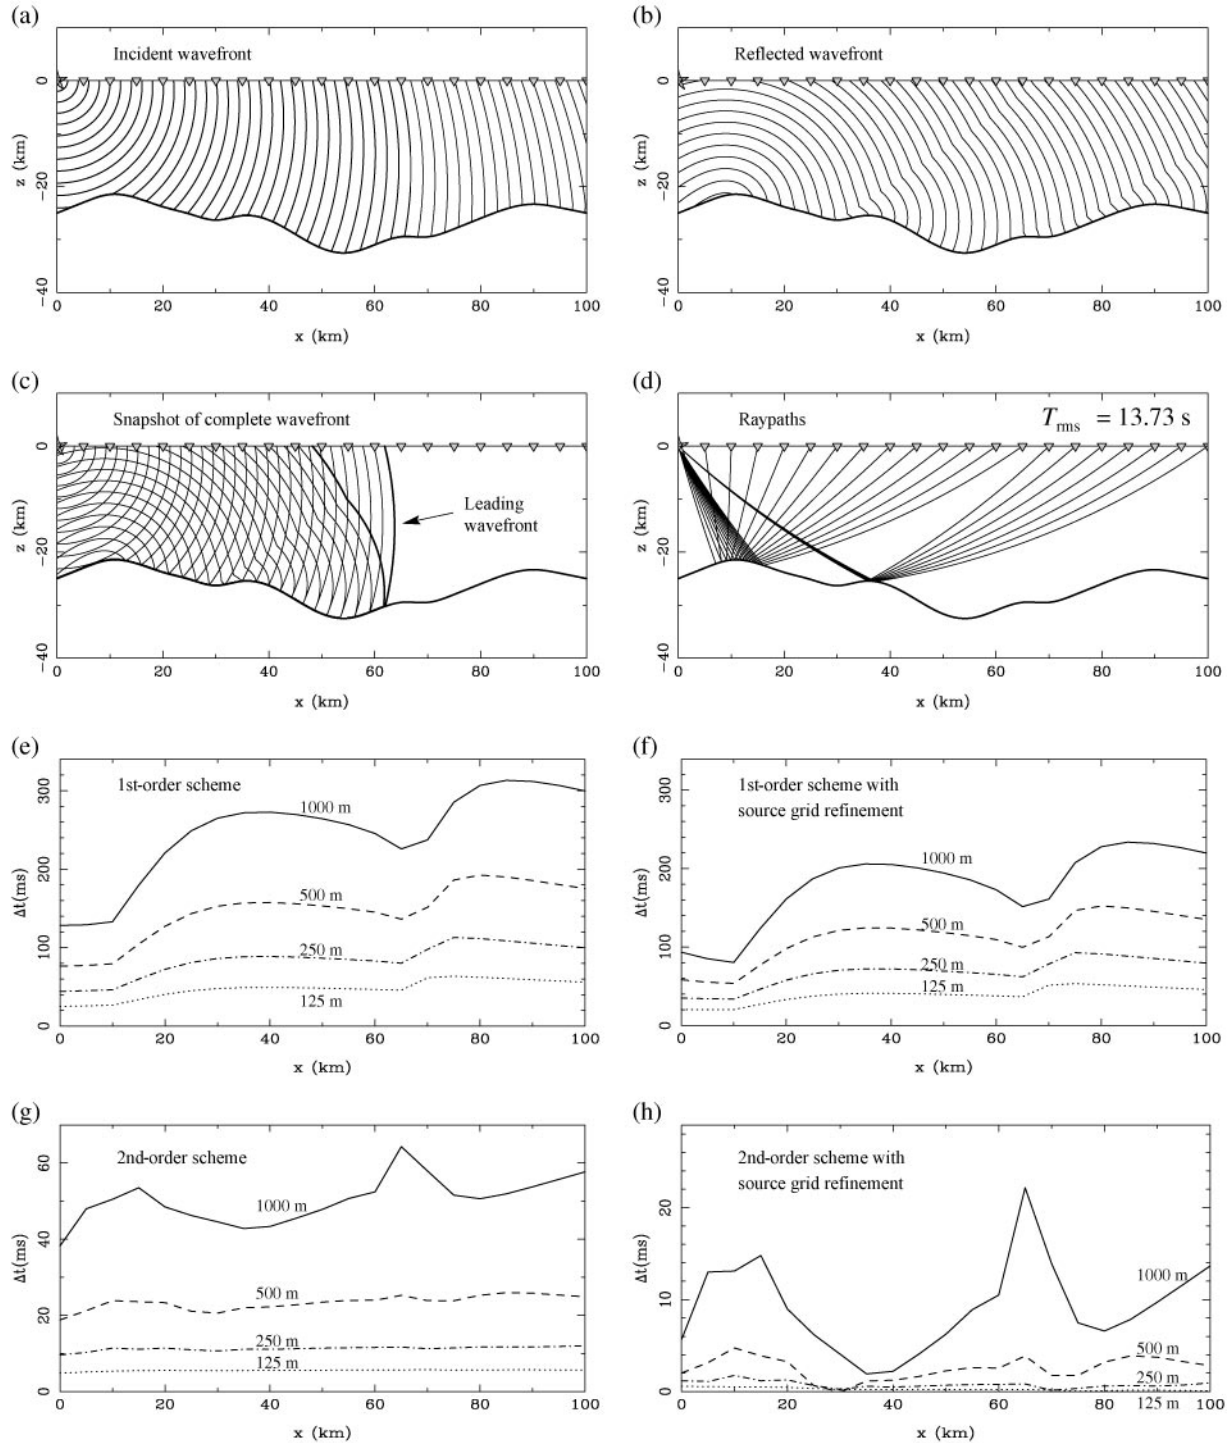
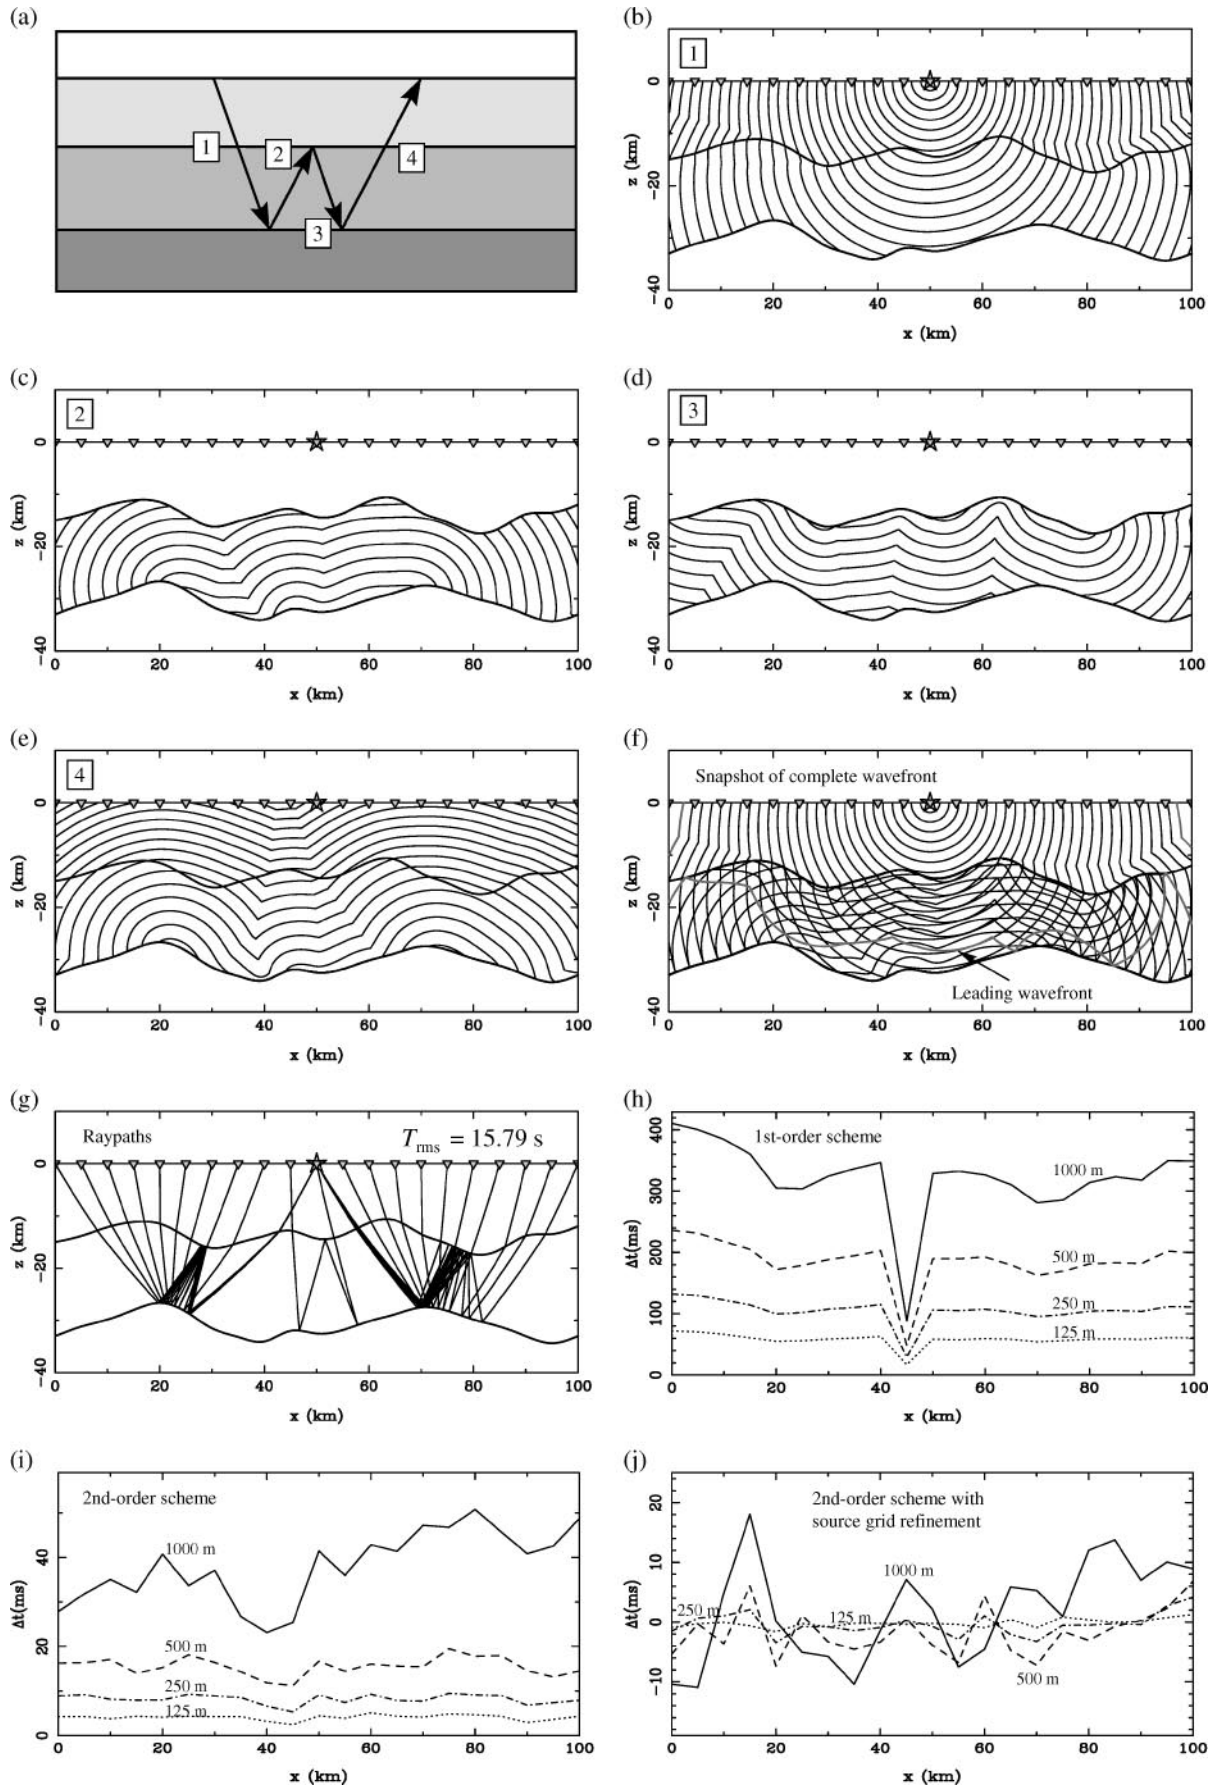


Figure 6. Tracking a simple reflection through a medium in which velocity varies linearly with depth. Wavefronts are plotted at 0.4-s intervals in all cases. (a) Incident wavefront emanating from the source; (b) reflected wavefront; (c) evolution of the complete multivalued wavefront (thick line denotes the leading wavefront); (d) raypaths (to 21 receivers) corresponding to the reflected phase. The rms traveltime to all receivers is given by T_{rms} ; (e)-(h) traveltime accuracy of four different schemes using four different grid sizes (1000 m, 500 m, 250 m, 125 m).



a coarser grid away from the source; therefore, it is similar to the method advocated by Kim and Cook (1999), who use an essentially nonoscillatory (ENO) scheme to solve the eikonal equation. The introduction of local grid refinement means that the computational front must at some stage pass from a finely spaced grid to a coarsely spaced grid and possibly back again. The latter case can produce numerical instabilities if simple interpolation is used (see Sethian, 1999). In the case of ENO schemes, the computational front is an expanding box; so, provided a rectangular grid of refined nodes is used, this situation will not arise. However, in the case of FMM, the computational front approximates the shape of the first-arrival wavefront. Unless the boundary of the refined grid has the same geometry, the wavefront may need to be upsampled at some stage.

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 5). As a result, the true edge of the refined grid conforms 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 the stability of FMM. In the following examples, the refined grid extends 10 coarse node points horizontally and vertically outward from the source and has a fivefold decrease in node spacing. For example, a velocity grid with a 1-km spacing has a refined grid with a 200-m node spacing surrounding the source to a distance of 10 km horizontally and vertically. Although this choice of refined grid appears ad hoc, we found that increasing the area of the refined grid and decreasing the node spacing has only a marginal effect on the accuracy of computed traveltimes.

EXAMPLES

Application of FMM requires that the velocity field and interfaces be discretely sampled; the node spacing chosen, along with the order of the upwind difference operator used, determines the accuracy of the calculated traveltimes. In our examples, the velocity field and interfaces are described by cubic B-spline functions in parametric form (Bartels et al., 1987), which are controlled by separate sets of velocity and interface vertices. Discrete sampling of these functions to any desired resolution is easily accomplished. Following, we present several examples of varying complexity to examine the behavior of the FMM scheme presented above.

The first example comprises a single reflection from an undulating interface that lies within a velocity field with a constant vertical gradient of 0.04 s^{-1} (see Figure 6). To track this phase, FMM is invoked twice: once for the incident traveltime field (Figure 6a) and once for the reflected traveltime field (Figure 6b). The geometry of the complete wavefront as it propagates through the medium can be obtained by stitching together matching isochrons from the incident and reflected

traveltime fields (Figure 6c). Source–receiver raypaths are computed a posteriori by following the traveltime gradient from the receiver back through the traveltime fields to the source (Figure 6d).

The velocity medium in Figure 6 is too complex to facilitate analytic solutions, so we use the ray-tracing method of Rawlinson et al. (2001) to examine the accuracy of the FMM scheme. The ray tracer is designed to find first-arrival reflected and refracted paths in layered media where the velocity varies linearly with depth and layer boundaries are described by cubic B-splines; this permits direct comparison with our FMM solutions. The desirable properties of the ray-tracing method are that it uses analytic solutions to track raypaths within a layer and an iterative Newton scheme to locate the intersection point of the raypath and the interface. The accuracy of the intersection point can be controlled, so highly accurate solutions are possible. 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 result in ray-tracing traveltimes that, for our purposes, can be considered exact.

Figures 6e–6h show error plots for four different FMM schemes with velocity grid spacings of 1000, 500, 250, and 125 m (see also Table 1). The error estimate Δt is equal to the difference between the FMM solution and the ray-tracing solution at each of the 21 receivers. The first-order scheme (Figure 6e) exhibits first-order convergence; including local grid refinement about the source (Figure 6f) improves the accuracy, but the size of the improvement suggests that errors introduced in the vicinity of the source are not principally responsible for the overall error. The second-order scheme (Figure 6g) is much more accurate than the first-order scheme (both with and without grid refinement) and converges more rapidly as the node spacing is reduced. When grid refinement about the source is used together with the second-order scheme (Figure 6h), the improvement in accuracy is again significant, and the convergence becomes near second order.

The second example (Figure 7) shows a multiple tracked through a structure composed of three layers. As in the Figure 6 example, velocity varies linearly with depth within each layer, but now we are trying to track four branches comprising three reflection events and two transmission events, as shown schematically in Figure 7a. Consequently, FMM is invoked four times to track the first-arrival wavefront corresponding to each branch (see Figures 7b–7e). A snapshot of the complete wavefront is shown in Figure 7f, and source–receiver raypaths are shown in Figure 7g.

The accuracy of the multistage FMM cannot be analyzed using the ray-tracing scheme in this case because it is not designed to compute multiples. Instead, we appeal to the stability of FMM and use traveltimes computed on a very fine grid as a proxy for the exact solution. In Figures 7h–7j, error estimates for the source–receiver traveltimes are plotted for four



Figure 7. Tracking a multiple through a layered medium in which velocity varies linearly with depth. Wavefronts are plotted at 0.4-s intervals in all cases. (a) Schematic illustration showing the four branches of the multiple that comprise the complete phase; (b) wavefronts corresponding to branch 1 (these emanate from the source); (c) wavefronts corresponding to branch 2; (d) wavefronts corresponding to branch 3; (e) wavefronts corresponding to branch 4; (f) evolution of the complete multivalued wavefront (thick line denotes the leading wavefront); (g) raypaths (to 21 receivers) corresponding to the fourfold multiple. The rms traveltime of the wavefront to all 21 receivers is given by T_{rms} ; (h)–(j) estimates of traveltime accuracy of three different schemes using four different grid sizes (1000, 500, 250, and 125 m).

different grid spacings (see also Table 2). The reference solution is obtained by applying the second-order FMM scheme on a 50-m grid with local grid refinement about the source. It is clear from these plots and Table 2 that the convergence behavior and accuracy of the different schemes are similar to the Figure 6 example. The only exception is the second-order scheme with grid refinement; the convergence appears to be approximately first order.

The multistage FMM scheme can track any number of phases in a layered medium. Figure 8 shows 10 different multipaths tracked between a single source and receiver through the structure used in Figure 7. The most complex of these paths comprise 18 reflection branches and two transmission branches. Two of the paths are pure transmissions: the direct arrival and a refraction phase that turns in the second layer. The latter phase was tracked using a three-stage FMM.

The next example is designed to test the robustness of the multistage FMM scheme by tracking a tenfold phase (Figure 9a) through a highly heterogeneous layered medium with lateral velocity contrasts as great as 8:1. The complexity of the evolving wavefront can be observed in Figure 9b, which clearly illustrates the multivalued nature of the wavefield. The corresponding raypaths to 21 receivers located on the surface (Figure 9c) also testify to the extreme heterogeneity of the velocity and interface structure. Since FMM is restricted to locating only first arrivals within a computational domain, the two-point ray paths favor the fast areas of the model.

Although the variations in wavespeed and interface structure are greater than we would expect to encounter in the earth, Figure 9d shows that traveltimes predicted by the first-order scheme to the 21 receivers appear to converge without instability as the grid size is reduced (see also Table 3). As in Figure 7, our proxy for the exact solution is obtained by using a second-order FMM with local grid refinement on a 50-m grid. The interesting feature of the second-order scheme (Figure 9e)

is that it converges more quickly than the first-order scheme; but when local grid refinement is applied, there is little improvement in accuracy.

Figure 10 shows a collection of 10 multiple arrivals tracked through the structure shown in Figure 9 between a single source and receiver. The most complex of these phases comprises 20 reflections and 22 transmissions. Compared to Figure 8, the paths are more clustered, a feature which may be attributed to the highly heterogeneous nature of the velocity distribution.

Our final example (Figure 11) shows a wavefront tracked through a medium similar to that used in Figure 9, except that the middle layer pinches out at both ends of the model. The phase in this case involves a reverberation between the bottom two interfaces; the wave bounces back and forth 19 times before returning to the surface. The complex nature of the wavefront that evolves from this process is clearly shown in Figure 11. Although the variations in wavespeed and interface




Figure 8. Ten different phases tracked between a source and receiver using the multistage FMM. The velocity and interface model is the same as that used in Figure 7.

Table 2. Summary of speed and accuracy tests for the reflection multiple example shown in Figure 7. CPU times are for a SunBlade 150 and include the time taken to read in the velocity field and interface surfaces.

Grid spacing (m)	Number of gridpoints	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	328.6	257.1	38.9	8.5	0.1	0.2	0.1	0.2
500	16 281	189.5	153.8	15.7	4.3	0.5	0.5	0.5	0.6
250	64 561	107.3	89.4	8.3	1.9	1.7	1.8	2.0	2.1
125	257 121	59.2	50.3	4.2	0.7	7.6	7.7	8.6	8.8

O = Order. R = Refined grid about source.

Table 3. Summary of speed and accuracy tests for the complex phase example shown in Figure 9. CPU times are for a SunBlade 150 and include the time taken to read in the velocity field and interface surfaces.

Grid spacing (m)	Number of gridpoints	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	460.7	430.4	130.8	118.0	0.2	0.3	0.2	0.3
500	16 281	255.2	249.7	40.2	36.0	0.6	0.7	0.7	0.8
250	64 561	148.1	145.5	13.9	11.5	2.4	2.5	2.7	2.8
125	257 121	87.1	85.7	4.3	3.5	9.9	10.1	11.4	11.6

O = Order. R = Refined grid about source.

structure exhibited by the model are extreme and the amplitude of the phase would be small, this example illustrates that the multistage FMM will work in any complex model.

DISCUSSION

The previous suite of examples shows that the multistage FMM preserves the computational speed and robustness of the original single-stage FMM and at the same time tracks a complete class of later arriving phases in layered media. The stability of the scheme can be observed even in the relatively simple case of Figure 7, where a fourfold multiple is constructed. In Figure 7d, a wavefront discontinuity at about $x = 60$ km im-

pings upon the deepest interface; Figure 7e shows this discontinuity preserved in the reflected wavefront. The stable propagation of wavefront discontinuities has traditionally been a major challenge for eikonal methods. Figure 7h shows that the traveltimes computed using the first-order scheme converge as the grid spacing is reduced, despite the complex nature of the wavefield. This convergence can even be observed in the somewhat pathological example of Figure 9, where a tenfold phase is tracked through a medium containing severe velocity and interface variations. The second-order scheme also appears to be stable in complex media.

The summary of speed and accuracy tests contained in Tables 1–3 shows that the multistage FMM is a computationally

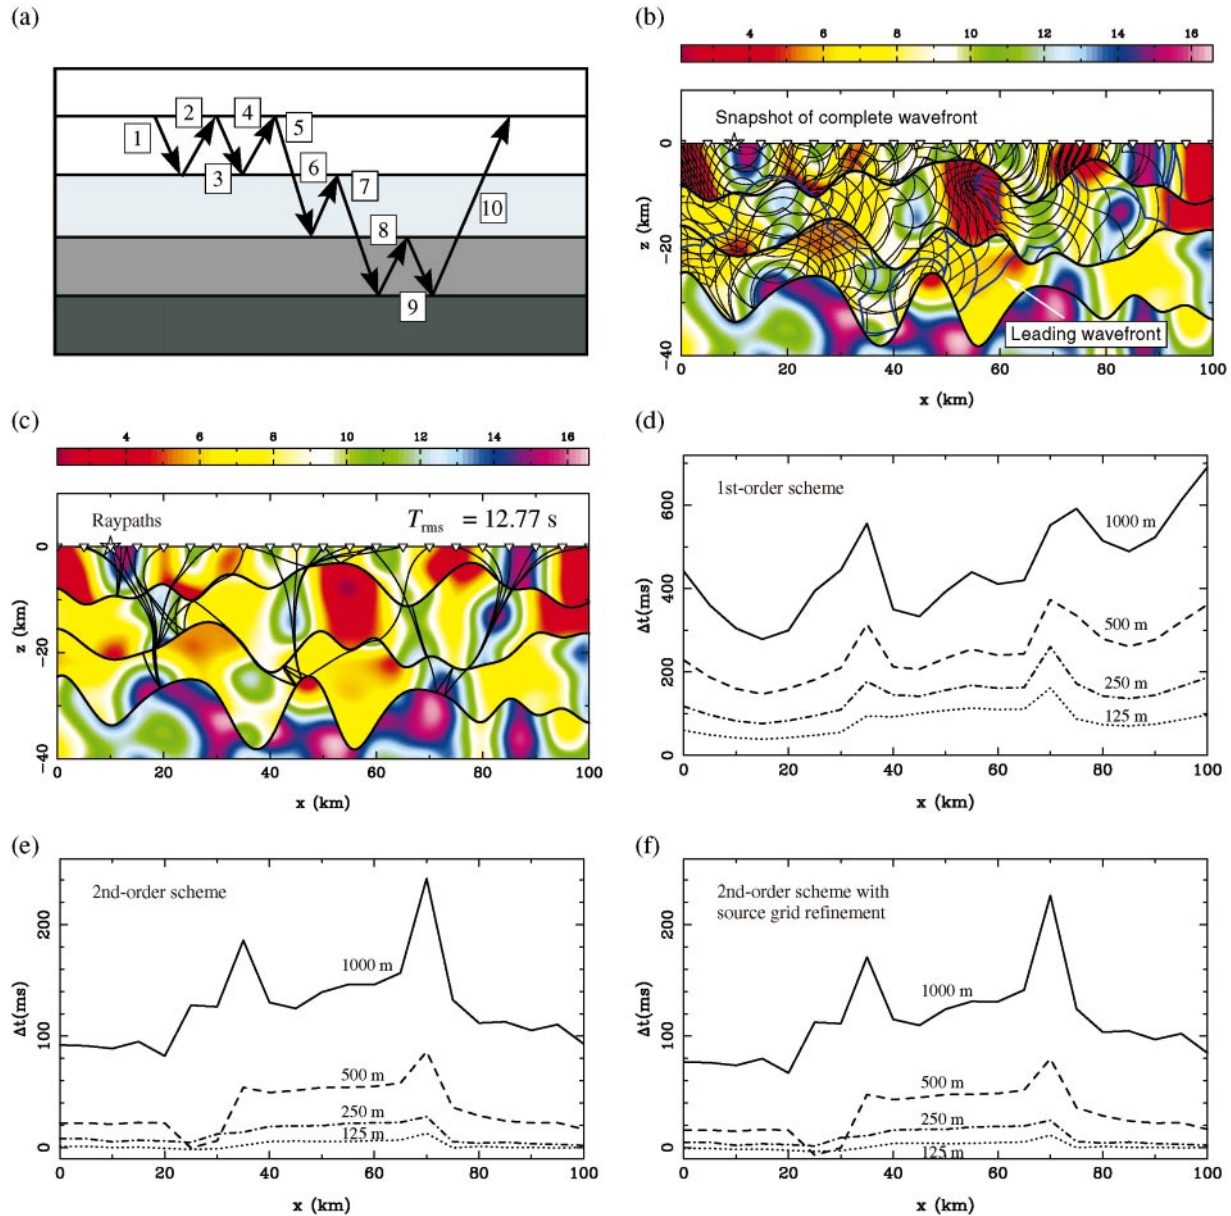


Figure 9. Tracking a tenfold multiple through a layered medium with extreme velocity and interface variations. (a) The 10 branches of the multiple that comprise the complete phase. (b) Evolution of the complete multivalued wavefield. Wavefronts (thin lines) are contoured at 0.4-s intervals; thick line shows the leading wavefront. (c) Raypaths to 21 receivers corresponding to the tenfold multiple. The rms traveltime of the wavefront to all 21 receivers is given by T_{rms} . (d)–(f) Estimates of traveltime accuracy for three different schemes using four different grid sizes (1000, 500, 250, and 125 m).

practical method for determining traveltimes, even if the required phase contains many reflection and refraction branches. It is also clear that a second-order scheme results in a significant increase in accuracy over a first-order scheme with little additional computational cost, and local grid refinement about the source can provide additional gains in accuracy. For instance, with a 125-m grid spacing in Figure 6 (Table 1), the second-order scheme with grid refinement is two orders of magnitude more accurate than the first-order scheme but requires only 15% more CPU time. The second-order scheme without grid refinement is an order of magnitude more accurate than the first-order scheme (at 125-m grid spacing) and takes 13% more CPU time.

As the complexity of the medium increases, the relative improvement of the second-order scheme compared to the first-order scheme appears to be approximately constant. However, the added benefit of local grid refinement about the source becomes less pronounced. For the example shown in Figure 6 (Table 1), the second-order scheme on a 125-m grid has an estimated error of 5.5 ms compared to 0.3 ms when local grid refinement is introduced. For the Figure 9 example, which is much more complex, the equivalent values are 4.3 and 3.5 ms, a much smaller difference. An explanation for this change is that the raypaths in Figure 9c emanate from the source at near-vertical




Figure 10. Ten different phases tracked between a source and receiver using the multistage FMM. The velocity and interface model is the same as that used in Figure 9.




Figure 11. A snapshot of a twentyfold reflection phase through a complex model that includes a layer pinch-out. The multi-valued wavefield is represented by thin lines plotted at 0.4-s intervals. The thick line denotes the leading wavefront.

trajectories, which is almost parallel to the vertical grid lines, a direction in which source error is a minimum (see Sethian and Popovici, 1999). Therefore, Grid refinement has little effect in decreasing the overall error. The fact that the second-order method converges more rapidly in Figure 9 than in Figure 6 also suggests that errors generated in the source vicinity are not as dominant in the former case.

A comparison of Tables 1 and 3 shows that increasing the heterogeneity of a velocity medium will decrease the accuracy of FMM. For example, at 1 km grid spacing, the second-order scheme with grid refinement has an estimated error of 50.6 ms (0.36%) for Figure 6 and 118.0 ms (0.93%) for Figure 9; at 125 m grid spacing, this reduces to 0.3 ms (0.002%) and 3.5 ms (0.027%), respectively. The velocity structure and phase shown in Figure 9 are much more complex than one would expect to encounter in the earth, so the level of error summarized in Table 3 is unlikely to be exceeded in practical applications. The multistage FMM compares favorably with other grid-based methods in terms of speed and accuracy (see Leidenfrost et al., 1999), and the added attraction of stability (unconditional in the first-order case) makes it a desirable scheme for complex media.

A limitation of the multistage FMM that we have alluded to before is that within a computational domain (i.e., one or more layers), only the first arrival is tracked. This can be seen clearly in Figures 7b–7e; the observed wavefront discontinuities mark the points at which later arriving information is lost. If all arrivals were tracked, then we would expect to see the formation of triplications (swallowtails), which would eventually lead to a significantly greater number of arrivals. Phases involving one or more consecutive bounces from the same interface also cannot be found using the multistage FMM; self-intersecting wavefronts must be tracked simultaneously if these arrivals are to be found.

The entropy-driven nature of FMM means that our new scheme only solves part of the complete multiarrival problem in layered media. However, robust grid-based methods can be tuned to work effectively in layered media, where reflections and transmissions form an integral part of the wavefield. It should therefore be possible to adapt other grid-based wavefront trackers—even those that are capable of constructing multiarrivals in continuous media—to work in layered media using our partitioning approach. To date, several papers develop the idea of exploiting phase space to find multivalued solutions of partial differential equations that describe high-frequency wave propagation in continuous media (Engquist et al., 2002; Fomel and Sethian, 2002; Osher et al., 2002). Our scheme for layered media would benefit from implementing these generalized techniques but at a greatly increased computational cost.

An alternative means of tracking multiple arrivals within a layer which is more consistent with the multistage FMM paradigm is to partition the multivalued 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 principal advantage of a partitioning approach is that first-arrival schemes are used to track wavefronts, so rapid computation is achieved. Combining these ideas with the wavefront reinitialization scheme in layered media would result in a powerful wavefront tracking tool.

As it currently stands, the multistage FMM can be practically applied to a number of problems encountered in observational seismology. For example, it can be used in coincident reflection or wide-angle traveltimes tomography; in earthquake relocation at a local scale or in local earthquake tomography; and in reflection processing to help suppress multiples and remove the effects of reverberations. The determination of amplitudes in addition to traveltimes would benefit these applications because it would help identify multiple arrivals which carry significant energy. One way of doing this is to solve the transport equation using the computed traveltimes field (e.g., Qian and Symes 2002); however, third-order accurate traveltimes are required to obtain first-order accurate amplitudes.

CONCLUSIONS

We have presented a multistage FMM for tracking multiple phases in complex layered media. The desirable properties of the new method include 1) stability (unconditional for first-order scheme), 2) computational speed and accuracy, and 3) the ability to track arrivals composed of any number of reflection and refraction branches. Application of the new scheme to highly heterogeneous media comprising severe variations in wavespeed and interface curvature supports these conclusions. We have also shown that the use of local grid refinement in the vicinity of the source can largely remove errors caused by a known singularity in the traveltimes field. Numerical results suggest that an order of magnitude improvement in traveltimes accuracy and near-quadratic convergence with a second-order scheme are possible. Although later arrivals (such as swallow-tails) caused by continuous variations in wavespeed cannot be tracked with the multistage FMM, interface generated phases form a major class of multiple arrival; as such, the new method may be applied in a wide range of areas including wide-angle tomography and the processing of coincident reflection data. Future work includes extension of the method to three dimensions, incorporating amplitudes, and tracking multiarrivals within a single layer.

ACKNOWLEDGMENTS

All figures were created using PGPLOT and xfig freeware graphics packages. This work was supported by Australian Research Council Discovery Projects DP0208039 and DP0451133.

REFERENCES

- Afnimar, and K. Koketsu, 2000, Finite difference traveltimes calculation for head waves travelling along an irregular interface: *Geophysical Journal International*, **143**, 729–734.
- Alkhalifah, T., and S. Fomel, 2001, Implementing the fast marching eikonal solver: Spherical versus Cartesian coordinates: *Geophysical Prospecting*, **49**, 165–178.
- Bartels, R. H., J. C. Beatty, and B. A. Barsky, 1987, An introduction to splines for use in computer graphics and geometric modelling: Morgan Kaufmann.
- Benamou, J. D., 1999, Direct computation of multivalued phase space solutions for Hamilton-Jacobi equations: *Communications on Pure and Applied Mathematics*, **52**, 1443–1475.
- Bevc, D., 1997, Imaging complex structures with semirecursive Kirchhoff migration: *Geophysics*, **62**, 577–588.
- Bulant, P., 1996, Two-point ray tracing in 3-D: *Pure and Applied Geophysics*, **148**, 421–447.
- Červený, V., 1987, Ray tracing algorithms in three-dimensional laterally varying layered structures, *in* Nolet, G., Ed., *Seismic tomography: With applications in global seismology and exploration geophysics*: D. Reidel Publishing Company, 99–133.
- Cheng, N., and L. House, 1996, Minimum traveltimes calculations in 3-D graph theory: *Geophysics*, **61**, 1895–1898.
- Chopp, D. L., 2001, Some improvements of the fast marching method: *SIAM Journal of Sci. Comput.*, **23**, 230–244.
- Engquist, B., O. Runborg, and A.-K. Tornberg, 2002, High-frequency wave propagation by the segment projection method: *Journal of Computational Physics*, **178**, 373–390.
- Farra, V., and R. Madariaga, 1988, Non-linear reflection tomography: *Geophysical Journal*, **95**, 135–147.
- Fischer, R., and J. M. Lees, 1993, Shortest path ray tracing with sparse graphs: *Geophysics*, **58**, 987–996.
- Fomel, S., and J. A. Sethian, 2002, Fast-phase space computation of multiple arrivals: *Proceedings of the National Academy of Science*, **99**, 7329–7334.
- Grechka, V. Y., and G. A. McMechan, 1996, 3-D two-point ray tracing for heterogeneous, weakly transversely isotropic media: *Geophysics*, **61**, 1883–1894.
- Hole, J. A., and B. C. Zelt, 1995, 3-D finite-difference reflection travel times: *Geophysical Journal International*, **121**, 427–434.
- Julian, B. R., and D. Gubbins, 1977, Three-dimensional seismic ray tracing: *Journal of Geophysics*, **43**, 95–113.
- Kim, S., 2002, 3-D eikonal solvers: First-arrival traveltimes: *Geophysics*, **67**, 1225–1231.
- Kim, S., and R. Cook, 1999, 3D traveltimes computation using second-order ENO scheme: *Geophysics*, **64**, 1867–1876.
- Kim, W., and C.-E. Baag, 2002, Rapid and accurate two-point ray tracing based on a quadratic equation of takeoff angle in layered media with constant or linearly varying velocity function: *Bulletin of the Seismological Society of America*, **92**, 2251–2263.
- Leidenfrost, A., N. Ettrich, D. Gajewski, and D. Kosloff, 1999, Comparison of six different methods for calculating traveltimes: *Geophysical Prospecting*, **47**, 269–297.
- Li, X. G., and T. J. Ulrych, 1993, Traveltimes computation in discrete heterogeneous layered media: *Journal of Seismic Exploration*, **2**, 305–318.
- Lutter, W. J., and R. L. Nowack, 1990, Inversion for crustal structure using reflections from the PASSCAL Ouachita experiment: *Journal of Geophysical Research*, **95**, 4633–4646.
- Moser, T. J., 1991, Shortest path calculation of seismic rays: *Geophysics*, **56**, 59–67.
- Nakanishi, I., and K. Yamaguchi, 1986, A numerical experiment on nonlinear image reconstruction from first-arrival times for two-dimensional island arc structure: *Journal of Physics of the Earth*, **34**, 195–201.
- Osher, S., L.-T. Cheng, M. Kang, H. Shim, and Y.-H. Tsai, 2002, Geometric optics in a phase-space-based level set and Eulerian framework: *Journal of Computational Physics*, **179**, 622–648.
- Pereyra, V., W. H. K. Lee, and H. B. Keller, 1980, Solving two-point seismic-ray tracing problems in a heterogeneous medium: *Bulletin of the Seismological Society of America*, **70**, 79–99.
- Podvin, P., and I. Lecomte, 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 J. A. Sethian, 2002, 3-D imaging using higher order fast marching traveltimes: *Geophysics*, **67**, 604–609.
- Prothero, W. A., W. J. Taylor, and J. A. Eickemeyer, 1988, A fast, two-point, three-dimensional ray-tracing algorithm using a simple step search method: *Bulletin of Seismological Society of America*, **78**, 1190–1198.
- Qian, J., and W. W. Symes, 2002, An adaptive finite-difference method for traveltimes and amplitudes: *Geophysics*, **67**, 167–176.
- Rawlinson, N., and M. Sambridge, 2004, Wavefront evolution in strongly heterogeneous layered media using the fast marching method: *Geophysical Journal International*, **156**, 631–647.
- Rawlinson, N., G. A. Houseman, and C. D. N. Collins, 2001, Inversion of seismic refraction and wide-angle reflection traveltimes for 3-D layered crustal structure: *Geophysical Journal International*, **145**, 381–401.
- Riahi, M. A., and C. Juhlin, 1994, 3-D interpretation of reflected arrival times by finite-difference techniques: *Geophysics*, **59**, 844–849.
- Sambridge, M. S., 1990, Non-linear arrival time inversion: Constraining velocity anomalies by seeking smooth models in 3-D: *Geophysical Journal International*, **102**, 653–677.
- Sethian, J. A., and A. M. Popovici, 1999, 3-D traveltimes computation using the fast marching method: *Geophysics*, **64**, 516–523.
- Sethian, J. A., 1996, A fast marching level set method for monotonically advancing fronts: *Proceedings of the National Academy of Science*, **93**, 1591–1595.
- , 1999, *Level set methods and fast marching methods*: Cambridge University Press.

- Symes, W. W., 1998, A slowness matching finite difference method for traveltimes beyond transmission caustics: 68th Annual International Meeting SEG, Expanded Abstracts, 1945–1948.
- Um, J., and C. Thurber, 1987, A fast algorithm for two-point seismic ray tracing: *Bulletin of the Seismological Society of America*, **77**, 972–986.
- van Trier, J., and W. W. Symes, 1991, Upwind finite-difference calculation of traveltimes: *Geophysics*, **56**, 812–821.
- Vidale, J. E., 1988, Finite-difference calculations of traveltimes: *Bulletin of the Seismological Society of America*, **78**, 2062–2076.
- 1990, Finite-difference calculations of traveltimes in three dimensions: *Geophysics*, **55**, 521–526.
- Zelt, C. A., and R. B. Smith, 1992, Seismic traveltime inversion for 2-D crustal velocity structure: *Geophysical Journal International*, **108**, 16–34.
- Zhao, Z., R. Kubota, S. Fumio, and S. Iizuka, 1997, Crustal structure in the southern Kanto-Tokai region derived from tomographic method for seismic explosion survey: *Journal of Geophysical Research*, **45**, 433–453.