Phase space methods for multi-arrival wavefronts

Juerg Hauser   Malcolm Sambridge   Nick Rawlinson

Key Words: Later arrivals, wavefront tracking, level set, Lagrangian, Marmousi model

ABSTRACT

Most body wave seismic imaging schemes only exploit information contained in the first arrival of a seismic record. Later arrivals in the wavetrain, however, contain additional structural information, as the corresponding rays tend to sample slower regions of a medium that are often avoided by first-arrival raypaths. Here we investigate a Lagrangian (ray-based) and an Eulerian (grid-based) approach for the calculation of later arrivals. The Eulerian approach is based on the level set method, which implicitly evolves a wavefront by solving a pair of PDEs over a gridded velocity field in phase space. Our Lagrangian solver also uses phase space and represents the wavefront by a set of points, which are progressively moved through the velocity field using local ray tracing, with linear interpolation used to maintain a constant density of points. We compare the two methods using a velocity model of the subduction zone in the Tonga region. In theory both approaches can provide traveltimes for later arrivals. Our results clearly show that the Lagrangian approach is currently superior to the Eulerian scheme for the prediction of multi-arrival traveltimes when computation speed, ease of implementation, and accuracy are considered. In our experiments the Lagrangian solver is up to 6000 times faster, and successfully predicts later arrivals for our source receiver configuration in the subduction zone example. We then demonstrate the robustness and efficiency of the Lagrangian solver by tracking later arrivals in a smoothed version of the Marmousi model. By placing source points in certain parts of this model, we are able to find more than 60 secondary arrivals at surface receivers.

INTRODUCTION

Both continuous and discontinuous variations in wavespeed can cause a seismic wave to travel to a receiver by more than one path, a phenomenon commonly referred to as multipathing. First-arrival raypaths tend to avoid slow anomalies; later arrivals do not avoid slow regions and hence contain additional information on seismic structure. Over the past few decades there has been a strong focus on the development of methods for the fast and reliable calculation of first-arrival traveltimes. These traveltimes have been computed either by ray tracing (e.g., Julian and Gubbins, 1977; Pereyra et al., 1980; Zelt and Ellis, 1988) or by wavefront tracking (e.g., Sun, 1992; Lambaré et al., 1992; Vinje et al., 1993). More recently implicit wavefront-tracking algorithms, such as those which use finite difference solutions of the eikonal equation, have also been employed (e.g., Vidale, 1988; Kim and Cook, 1999; Qian and Symes, 2002).

Research School of Earth Sciences
The Australian National University
Canberra ACT 0200 Australia
Phone: +61 2 6125 1520
Facsimile: +61 2 6257 2737
Email: juerg@rses.anu.edu.au

The traditional approach for the calculation of traveltimes has been ray tracing between a source and receiver, which can be achieved by shooting or bending rays (Julian and Gubbins, 1977). In the shooting method, rays are initiated at the source point with different initial directions and propagated through the medium by solving an initial value formulation of the ray-tracing equation. The initial trajectories of rays are then updated until they hit the receiver within a given tolerance. In the bending method an initial path between source and receiver is updated using a boundary value formulation of the ray equation until it corresponds to a stationary path (i.e., satisfies Fermat's Principle). However, the raypath calculated by ray tracing is not necessarily the first-arrival raypath.

Since the late 1980s there has been considerable interest in and development of numerous grid-based algorithms for the efficient calculation of arrival times using various finite-difference methods. One of the first attempts to compute the first-arrival traveltime field using such a technique was made in two dimensions by Vidale (1988), who later extended it to three dimensions (Vidale, 1990). The scheme involves progressively integrating the traveltimes along an expanding square in two dimensions or an expanding cube in three dimensions. In an alternative approach Moser (1991) computes seismic raypaths by tracking the shortest traveltime path through a predefined network, which links the nodes of a gridded velocity medium. In the shortest path method, all paths from one point are constructed simultaneously.

Essentially non-oscillatory schemes (ENO) (Harten et al., 1987; Shu and Osher, 1988; Shu and Osher, 1989), and weighted essentially non-oscillatory schemes (WENO) (Liu et al., 1994; Jiang and Shu, 1996; Jiang and Peng, 2000), are algorithms for solving Hamilton-Jacobi type equations, for example hyperbolic conservation laws. The eikonal equation has the form of such a conservation law, which means that first-arrival traveltimes can be computed using, for example, second and third order WENO and ENO schemes (e.g., Kim and Cook, 1999; Qian and Symes, 2002; Buske and Kaestner, 2004).

Another eikonal equation based technique for computing first-arrival traveltimes is the Fast Marching Method (e.g., Sethian and Popovici, 1999) or FMM, which is a grid-based wavefront evolution method capable of tracking first arrivals. FMM can be implemented with unconditional stability and has recently been extended to allow for the tracking of refracted and reflected wavefronts in layered media (Rawlinson and Sambridge, 2004). Important advantages of these finite-difference approaches, and in particular of FMM, compared to traditional ray tracing, are their computational efficiency, algorithmic simplicity, robustness, and solution completeness for first arrivals.

However, the fundamental limitation of these solvers is that they cannot track later arrivals of a wavefront within a continuous medium. There have been attempts to compute multivalued traveltime fields using only a first-arrival solver, but these approaches often include a rather ad hoc procedure for dividing the domain into subregions, followed by use of a first-arrival technique in each subregion. The solutions for

the different subregions are then superimposed to provide multivalued traveltimes for the whole computational domain (e.g., Fatemi et al., 1995; Benamou, 1999).

Lagrangian approaches to the problem of seismic wavefront tracking were introduced by Lambaré et al. (1992) and Vinje et al. (1993). Here a set of points is used to represent the wavefront, which is iteratively evolved by using local ray tracing with a given time step. After each iteration, the density of points is evaluated, and interpolation is used to add points (using a metric distance criterion) to regions of the wavefront that are undersampled. Sun (1992) recognised that the angular distance and therefore curvature should also be taken into consideration when adding points to the wavefront. Lambaré et al. (1996), in two dimensions, and Lucio et al. (1996), in three dimensions, used the Hamiltonian formulation of ray theory for their wavefront-tracking algorithm. In the Hamiltonian formulation (Chapman, 1985) the rays in normal space are replaced by bicharacteristics in phase space.

Osher and Sethian (1988) pioneered the field of interface tracking. They developed the level set method, which keeps track of an interface by expressing it as the zero level set (zero contour line) of a function representing the signed distance to the interface. The interface is then evolved implicitly as the signed distance function is updated on an Eulerian grid. The level set method cannot track self-intersecting interfaces and as a result cannot be used directly to track later-arriving wavefronts. Osher et al. (2002) therefore create a three-dimensional phase space from two-dimensional normal space by simply using the local normal direction to the wavefront as the third coordinate. In phase space, a self-intersecting wavefront unravels to become a non-self-intersecting line, which in the level set literature is known as the bicharacteristic strip. This strip can then be tracked in phase space and projected back into normal space to provide the multivalued geometrical optics solution of the eikonal equation. In short they describe the bicharacteristic strip in phase space as the intersection of the zero level sets of two three-dimensional signed distance functions. Computing the multivalued geometrical solution of the eikonal equation has always been a driving force in the level set community (e.g., Osher et al., 2002; Fomel and Sethian, 2002; Cockburn et al., 2005), and these attempts by applied mathematicians stand in stark contrast to decades of ray tracing in seismology. Level set proponents claim that their approach is superior to a Lagrangian approach for the computation of multi-arrival traveltimes due to the implicit representation of the wavefront. In these schemes, however, the spatial resolution of the wavefront is always controlled by the resolution of the underlying grid.

To date, no significant applications of these techniques have emerged in geophysics (Benamou, 2003). Wavefront tracking in phase space has also not been extensively investigated in seismology, particularly in applications outside the exploration field. In addition, many of the suggested algorithms have only been tested in relatively simple two- or three-dimensional media, that is, in media with smoothly varying velocities and restricted peak to peak amplitudes. In the following we will compare an Eulerian and a Lagrangian algorithm for the computation of traveltimes in two-dimensional media. First we introduce the concept of phase space and show how three-dimensional phase space can be constructed from two-dimensional normal space. After describing both solvers, we test their accuracy using a one-dimensional model with a constant velocity gradient and then apply them to a model of the subduction zone in the Tonga region. We then demonstrate the power of our Lagrangian solver by using it to track large numbers of arrivals in the strongly heterogeneous Marmousi model.

METHOD

Eulerian Scheme

Fast marching and level set methods are techniques for solving interface propagation problems. The basic theory has been developed and summarised in books by Sethian (1999) and Osher and Fedkiw (2003). In the level set method the interface is described as the zero level set (the zero contour line) of a higher-dimensional function (i.e., a signed distance function). The signed distance function is defined as the distance from a point of the underlying grid to the closest point on the interface, and is negative if the grid point is on one side of the interface and positive if the grid point is on the other side of the interface. The zero contour or zero level set of the signed distance function ϕ then corresponds to the original interface and can be written as

$$\phi(\mathbf{x}(t),t) = 0 , \tag{1}$$

where $\mathbf{x}(t)$ is the position of the interface (i.e., the zero level set) at the time t. The time derivative of (1) is found with the chain rule

$$\phi_t + \nabla\phi(\mathbf{x}(t),t) \cdot \mathbf{x}_t(t) = 0 , \tag{2}$$

where ϕ_t and $\mathbf{x}_t$ denote derivatives with respect to time. The scalar speed function F is defined as the speed in the outward normal direction:

$$F = \mathbf{x}_t(t) \cdot \mathbf{n} , \tag{3}$$

where

$$\mathbf{n} = \frac{\nabla\phi}{|\nabla\phi|} . \tag{4}$$

Hence the evolution equation for the signed distance function ϕ is:

$$\phi_t + F|\nabla\phi| = 0 \quad \text{given} \quad \phi(\mathbf{x},t=0) . \tag{5}$$

This is the level set equation given by Osher and Sethian (1988), and corresponds to equations describing advective transport in an incompressible fluid. From a conservation law point of view, equation (5) states that for a given point the change of a property in time (for example, the concentration of a substance in a fluid) is equal to the flux of this property in the direction of the gradient. The signed distance function is well behaved when the absolute value of the gradient is equal to one for every point of the computational domain. In this situation equation (5) reduces to $\phi_t = -F$, and the values of ϕ either increase or decrease, depending on the sign of F. When $F > 0$, the interface moves in the outward normal direction, and when $F < 0$ the interface moves in the inward normal direction. When $F = 0$ the equation reduces to $\phi_t = 0$, and hence ϕ is not updated and the interface does not move.

For nodes of the grid that are equidistant from at least two points on the front the gradient of the signed distance function is discontinuous. This means that any numerical method used for evolving the signed distance function has to show a reasonable behaviour at the occasional kinks where the gradient fails to exist. If ϕ is not a well-behaved smooth signed distance function, the level set equation (5) must be solved numerically by treating it as a Hamilton-Jacobi equation. The literature contains a large number of numerical schemes for solving hyperbolic conservation laws and especially Hamilton-Jacobi type equations (e.g., Liu et al., 1994; Jiang and Shu, 1996; Jiang and Peng, 2000). In these schemes an approximation to the gradient is used in cases where it breaks down, often based on the smoothest set of nodes.

The level set method has the advantage that there is no need to add or remove points from the interface as it evolves. Topological changes like breaking and merging of the evolving front are handled naturally. Geometrical properties of the interface like the normal to the front can be defined using the gradient of the signed distance function ϕ. The signed distance function also specifies on which side of the interface a point of the underlying grid lies.

For an external velocity field $\mathbf{v}$, the level set equation (5) can be written as

$$\phi_t + \mathbf{v} \cdot \nabla\phi = 0 \ , \tag{6}$$

where $F|\nabla\phi| = \mathbf{v} \cdot \nabla\phi$.

The level set method cannot describe a self-intersecting interface. This is because of the fundamental principle underlying the level set method, which requires that the minimum entropy or viscous solution is found. However, a wavefront passing through a low-velocity region may develop a self-intersecting swallowtail pattern. Osher et al. (2002) therefore unfolds the self-intersecting wavefront into phase space. For seismic wavefront propagation in two dimensions, each point (x, y) on the wavefront can be represented in a three-dimensional phase space by simply using the local normal direction θ as the third coordinate (Figure 1(a)). An example is shown in Figure 1(b), where the circular wavefront at the base of the cube becomes the three-dimensional spiral strip, which is periodic in the θ direction and referred to as the bicharacteristic strip. The velocity field for the motion in phase space is given by the following equation, where c is the scalar velocity in normal space, c_x the derivative of the velocity in the x direction and c_y the derivative in the y direction:

$$\mathbf{v} = \begin{pmatrix} c \cdot \cos\theta \\ c \cdot \sin\theta \\ c_x \cdot \sin\theta - c_y \cdot \cos\theta \end{pmatrix} . \tag{7}$$

Equation (7) can also be viewed as an initial value formulation of the ray-tracing equation (cf. equation (8)). We represent the velocity field using B-splines because equation (7) contains the derivatives of the velocity field and a B-spline representation provides smooth first derivatives.

In general the level set method can be formulated for the evolution of an m-1 dimensional manifold in an m-dimensional space. This means we can evolve a line in two dimensions or a surface in three dimensions. Therefore Osher et al. (2002) describe the wavefront in phase space (i.e., the bicharacteristic strip) as the intersection of two implicit surfaces (Figure 1(c)). The two implicit surfaces are then independently evolved using the level set method. We solve the level set equation using the fifth-order WENO scheme described by Jiang and Peng (2000). When the wavefront for a given moment in time needs to be extracted, the intersection of the two surfaces is computed and mapped back into normal space.

For a point source the bicharacteristic strip in phase space is given by a line with constant x and y, and θ varies between $-\pi$ and $+\pi$. The two surfaces are then given by planes of constant x and constant y. The signed distance at a node of the underlying grid with position (x, y) to the surface are is defined by $x - x_s$ and $y - y_s$, with (x_s, y_s) being the source location.

The level set method can generally only represent features of a surface that extend over more than one grid cell (e.g., Qian and
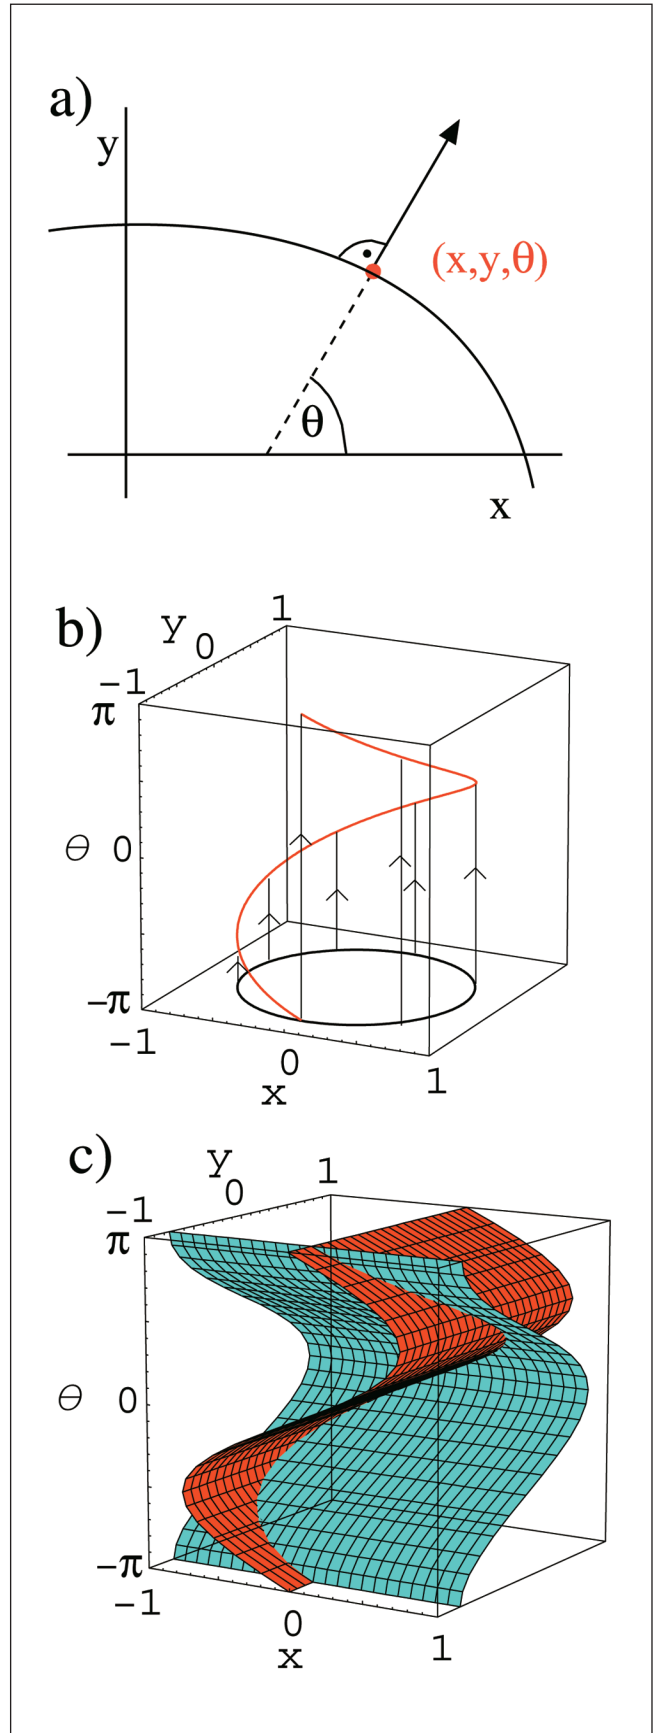

Fig. 1. Phase space representation of a two-dimensional wavefront. (a) Ray trajectory from the wavefront (local wavefront normal) supplies third dimension θ. (b) Wavefront in normal space (black line) and corresponding bicharacteristic strip (red line) in phase space. (c) Representing the bicharacteristic strip as the intersection of two surfaces.
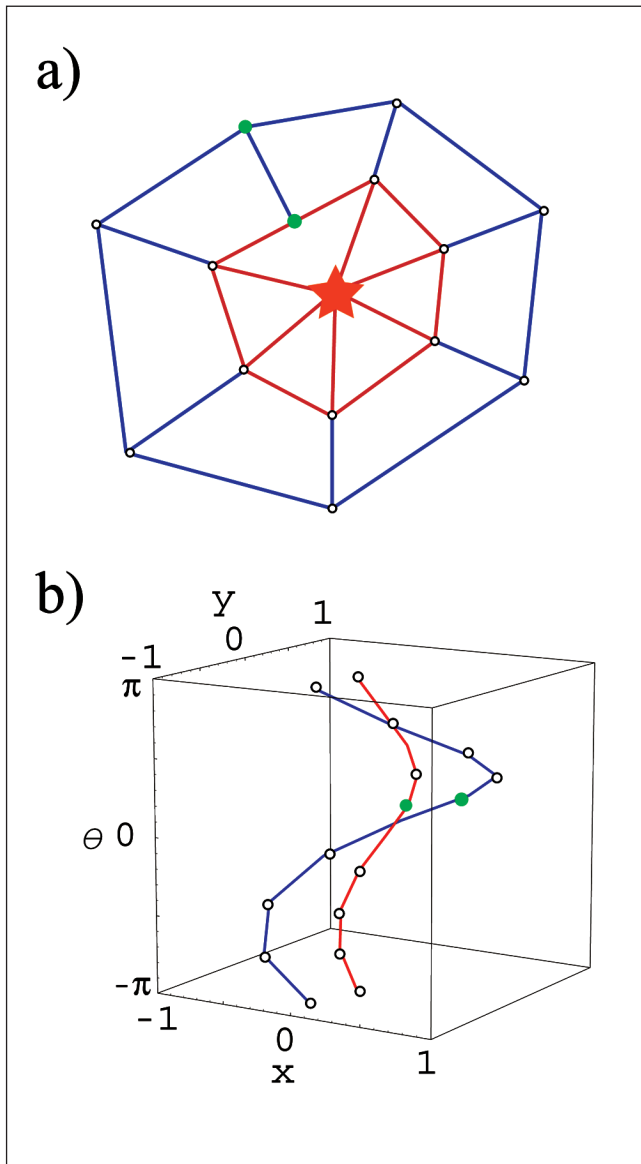
Fig. 2. Lagrangian scheme. (a) Wavefronts in normal space. (b) Corresponding bicharacteristic strip in phase space. A point (green) is inserted if the phase space distance between two neighbouring points is above a certain threshold.

Leung, 2004). However increasing the grid resolution by a factor of two increases the computation time by a factor of sixteen because of the requirement that the interface does not move more than one grid cell per time step. This inherent level of scaling means that the available memory and computation time will always limit the accuracy of the method.

Lagrangian Scheme

Within a Lagrangian framework, the bicharacteristic strip is represented by a finite set of points. As noted above, for a point source the bicharacteristic strip in phase space is given by a line with constant x and y, and θ between $-\pi$ and $+\pi$, where x and y are the source location. We then represent this line in phase space by a uniformly distributed set of points. The bicharacteristic strip is evolved through the medium in a set of discrete time steps. For a given time step δt, the strip is updated using a two-stage procedure. In the first stage all points are evolved in time by solving the initial value formulation of the ray-tracing equation (see Rawlinson and Sambridge, 2004) using a fourth-order Runge Kutta solver:
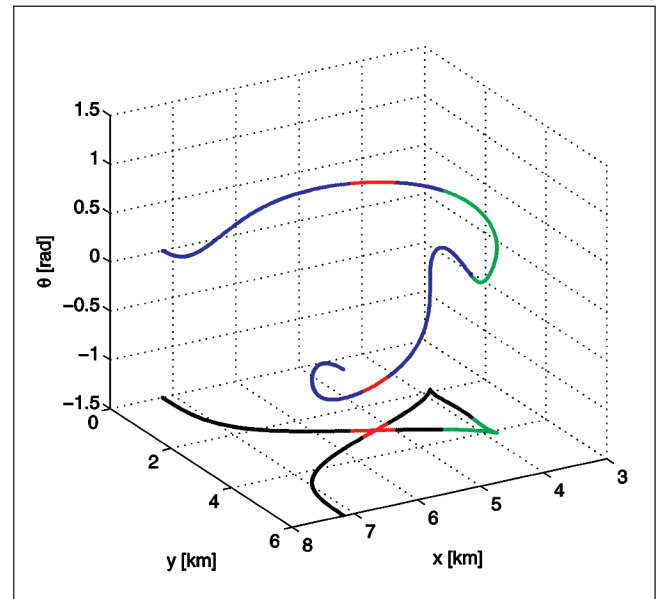

Fig. 3. Swallowtail pattern of a wavefront. The intersecting segments in normal space (red lines) do not intersect each other in phase space. The sharp corners in normal space (green segments) are given by a smooth representation in phase space.

$$\begin{pmatrix} \dfrac{dx}{dt} \\[6pt] \dfrac{dy}{dt} \\[6pt] \dfrac{d\theta}{dt} \end{pmatrix} = \begin{pmatrix} c \cdot \cos\theta \\ c \cdot \sin\theta \\ c_x \cdot \sin\theta - c_y \cdot \cos\theta \end{pmatrix}, \tag{8}$$

where c_x is the derivative of the velocity in the x direction, c_y the derivative in the y direction, and the position of the point in phase space is given by (x, y, θ). In the second stage, points are added or removed from the bicharacteristic strip, depending on their distance apart in phase space (Figure 2). The aim is to keep a fixed density of points along the bicharacteristic strip in order to minimise the loss of detail. Increasing the number of points on the bicharacteristic strip by a factor of two increases the computation time and memory requirements by a factor of two only. For the calculation of the distance between two neighbouring points we scale phase space so that the two metric coordinates x and y also lie in $[-\pi, +\pi]$. This scaling defines the norm distance measure in phase space and hence a definition of point density. A point is added if the spacing between two neighbouring points becomes larger than twice the initial distance, and a point is removed if the distance between the two neighbours of a given point is smaller than half the initial separation. We add points to the bicharacteristic strip using linear interpolation. Although higher order interpolation could be used, our results for the constant velocity gradient will show that the accuracy achieved by using a linear interpolation is already sufficient for practical applications.

Removing points can decrease accuracy but will invariably increase the efficiency of the scheme. In earlier studies, ray density has been defined only in normal space – for example, the metric distance between neighbouring rays (Lambaré et al., 1992; Vinje et al., 1993) or the angular distance (Sun, 1992). These definitions of ray density tend to encounter difficulties if the wavefront starts to develop a swallowtail pattern; they are only loosely correlated to the complexity of the ray field. The use of a phase space distance metric is a key element of our Lagrangian method and is superior to the alternative of using a metric defined in normal space. Lambaré

et al. (1996) use a similar criterion for the ray density in the phase space within their Hamiltonian formulation of ray theory.

Figure 3 illustrates how the bicharacteristic strip stays non-self-intersecting and piecewise smooth for a self-intersecting wavefront with sharp corners. The two intersecting segments of the wavefront are mapped to different positions in phase space because the directions of their local wavefront normals differ substantially from each other (red segments in Figure 3). If points are added to the wavefront in normal space near the sharp corner in the green region, a linear interpolation cannot provide adequate results. In phase space however a linear interpolation is simple to apply along the smooth bicharacteristic strip (green segment in Figure 3).

Extracting Arrival Time Information

So far we have shown how the Eulerian and the Lagrangian solver can be used to calculate the position of the wavefront for a given moment in time. However, in seismology we often want to know the traveltime, and the raypath to a receiver is more useful information. We extract this arrival-time information in a post-processing step once all the wavefronts have been calculated.

A point receiver in real space projects into a line in phase space. In the Eulerian approach the square of the distance to the closest point on the wavefront is given by the minimum of the sum of the squares of the two signed distance functions along the line in phase space representing the receiver. This information is stored for each time step and each receiver. The recorded distance forms a local minimum if a wavefront has passed a receiver. Fitting a parabola through the three points defining
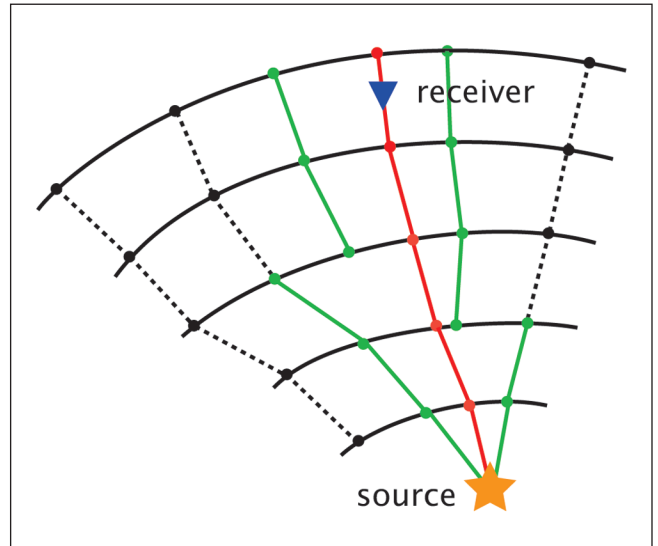

Fig. 4. Determination of a raypath using the wavefronts calculated by the Lagrangian solver. Given the green and black points on the wavefront and the corresponding raypath segments, a raypath (red line) for a receiver is interpolated back to the source between the green known neighbouring raypath segments.

the local minimum allows an arrival time to be interpolated. However two arrivals can only be distinguished from each other if their temporal separation is bigger than the size of the time step. The squared distance is also only measured for a limited number of points on the wavefront, and therefore not all the
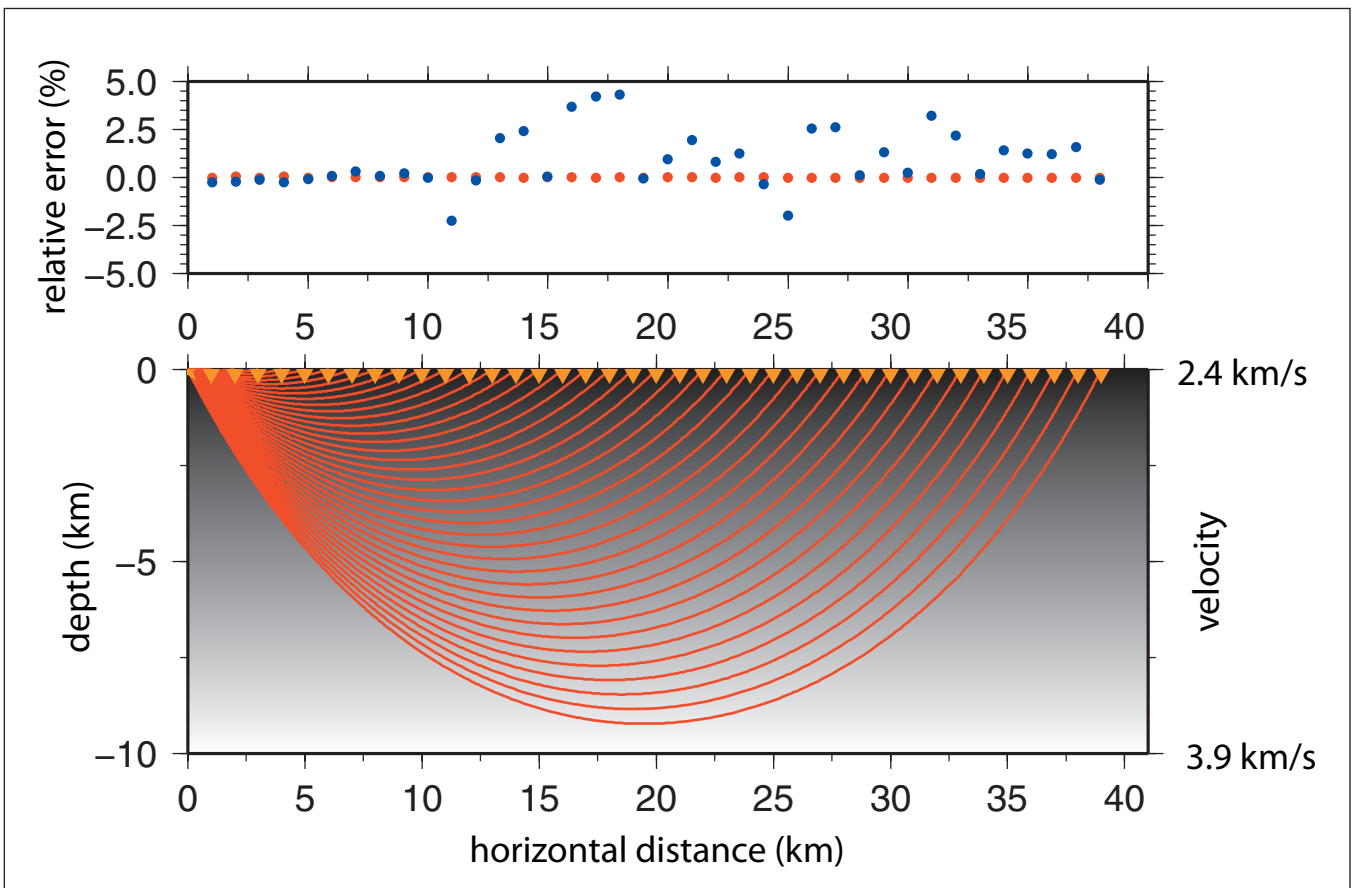

Fig. 5. Comparison of the accuracy of the Lagrangian and Eulerian solvers in a medium described by a constant velocity gradient. In the top figure the relative difference between the computed and analytical traveltime is plotted as a function of receiver position in blue for the Eulerian solver and red for the Lagrangian approach. The velocity model is plotted at the bottom and the velocity increases linearly with depth and is 2.4 km/s at the surface and 3.9 km/s at the bottom. The raypaths in the figure at the bottom were computed using the Lagrangian solver.

local minima of the squared distance are equal to zero, because even if a wavefront lies on a receiver the distance to the closest point might not be equal to zero. We therefore need to define a threshold in order to distinguish between true arrivals and the passage of endpoints of a swallowtail close to a receiver.

In the Lagrangian scheme the wavefront is explicitly defined at each time step. This allows us to identify the consecutive wavefronts and adjacent raypaths that together bound a receiver. The raypath segment and traveltime can then be computed using linear interpolation. This scheme is similar to an approach by Vinje et al. (1993).

Source-receiver raypaths can be estimated a posteriori in the case of the Lagrangian scheme by exploiting the connectivity between points on the wavefront for all time steps. Once wavefront tracking is finished, we have a set of raypaths (green and black lines in Figure 4) through the medium for all points of the bicharacteristic strip (black and green points in Figure 4). For the determination of a raypath between the source and a given receiver we begin at the receiver and progressively interpolate between known neighbouring segments back to the source. This approach is illustrated in Figure 4 where the red raypath is interpolated back to the source between the green known raypath segments. Having extracted the raypath, an estimate for the relative amplitude can be determined by solving the dynamic ray-tracing equations along the given path (e.g., Vinje et al., 1993).

EXAMPLES

The two algorithms are implemented in FORTRAN 95. All computations are performed on a PC with a 3.2 GHz Intel Pentium 4 processor and 2 Gb RAM, running GNU/Linux. First the accuracy of both approaches is tested, using a one-dimensional velocity model with a constant gradient. We then use a two-dimensional laterally heterogeneous seismic velocity model of a subduction zone in the Tonga region (Conder and Wiens, 2006) for the comparison between the Eulerian and Lagrangian solver. This test confirms the fact that the usefulness of the Eulerian solver is severely limited by the nature of its trade-off between grid resolution, accuracy, and computation time. We then demonstrate the capabilities of our Lagrangian solver using the Marmousi model from the 1996 INRIA workshop on multi-arrival traveltimes (see <http://www.caam.rice.edu/~benamou/traveltimes.html>).

Accuracy of First arrivals for a Constant Velocity Gradient

For a constant velocity gradient model, the analytical solution for the travel time is known (Sheriff and Geldart, 1995). The model for this test is given by a velocity of 2.4 km/s at the top and 3.9 km/s at the bottom with a constant gradient in between (Figure 5). While this model does not generate later arrivals it provides an estimate for the accuracy of the two methods. The Lagrangian solver uses 150 points to represent the point source in phase space; the error between the numerical and analytical solution is in this case less than 0.1% for all receivers. The errors in the Eulerian approach tend to be larger but still less than 5% for all receivers (Figure 5). The grids for the signed distance functions in the Eulerian approach consist of 161×41×51 nodes. For one time step the values for all nodes of the two grids representing the signed distance function have to be calculated; by contrast the Lagrangian approach only requires 150 points to be updated, although this increases as the wavefront expands. This explains why the Lagrangian approach takes 7 seconds of CPU time while the Eulerian approach requires 35 minutes of CPU time. More advanced level set methods might be more accurate and efficient but they would still require the solution
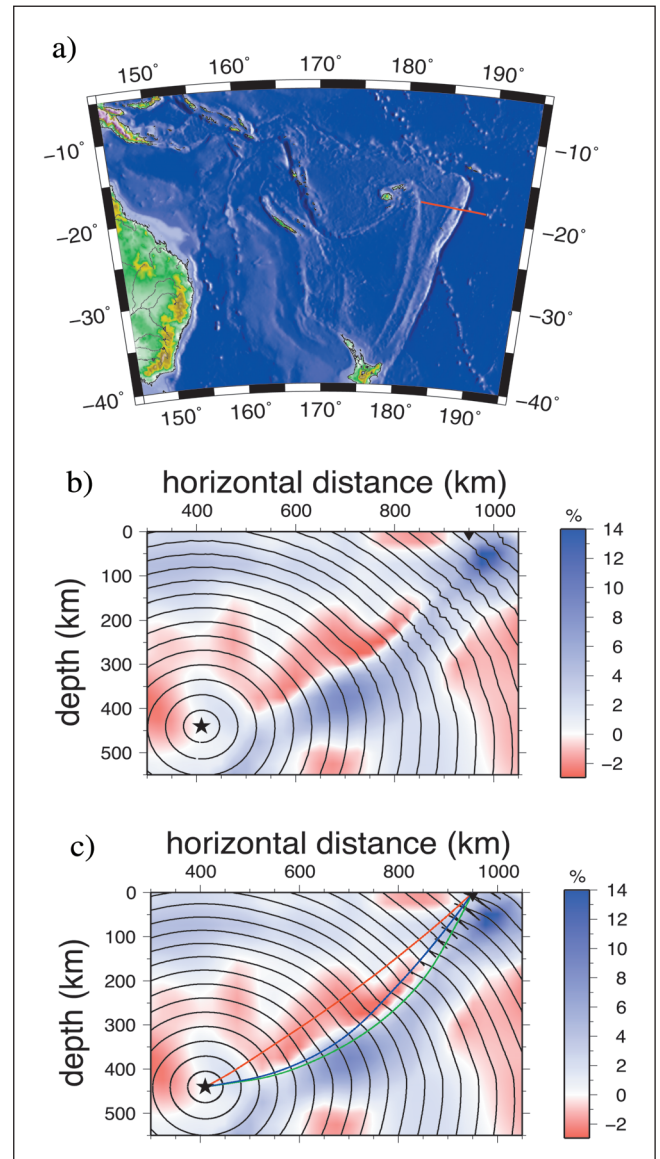

Fig. 6. Later arrivals in a subduction zone. (a) Orientation of the two-dimensional P-wave velocity model (Conder and Wiens, 2006). (b) Wavefronts calculated using the Eulerian solver. The velocity is plotted as perturbation with respect to a local one-dimensional model. (c) Wavefronts calculated using the Lagrangian scheme. We observe three arrivals at the receiver and their raypaths are plotted red for the first, green for the second, and blue for the third arrival.

of two three-dimensional problems in order to solve what is essentially a one-dimensional problem. The lower accuracy of the Eulerian solver could also be partly due to the limited accuracy of the algorithm used for the extraction of the arrival times.

Computing Multi-Arrivals in a Subduction Zone.

We use a two-dimensional P-wave velocity model of a subducting slab in the Tonga region (Conder and Wiens, 2006) to illustrate the basic features of the Eulerian and Lagrangian solver. We expect the appearance of multiple arrivals due to the high-velocity anomaly associated with the subduction of the cold oceanic lithosphere. The plane of the model is perpendicular to the slab, and the model extends for 1400 km in the east-west direction and 700 km in the depth direction. P-wave velocity values are given by Conder and Wiens (2006) for a Cartesian grid consisting of 57×29 nodes spaced 25 km horizontally and vertically (Figure 6(a)).
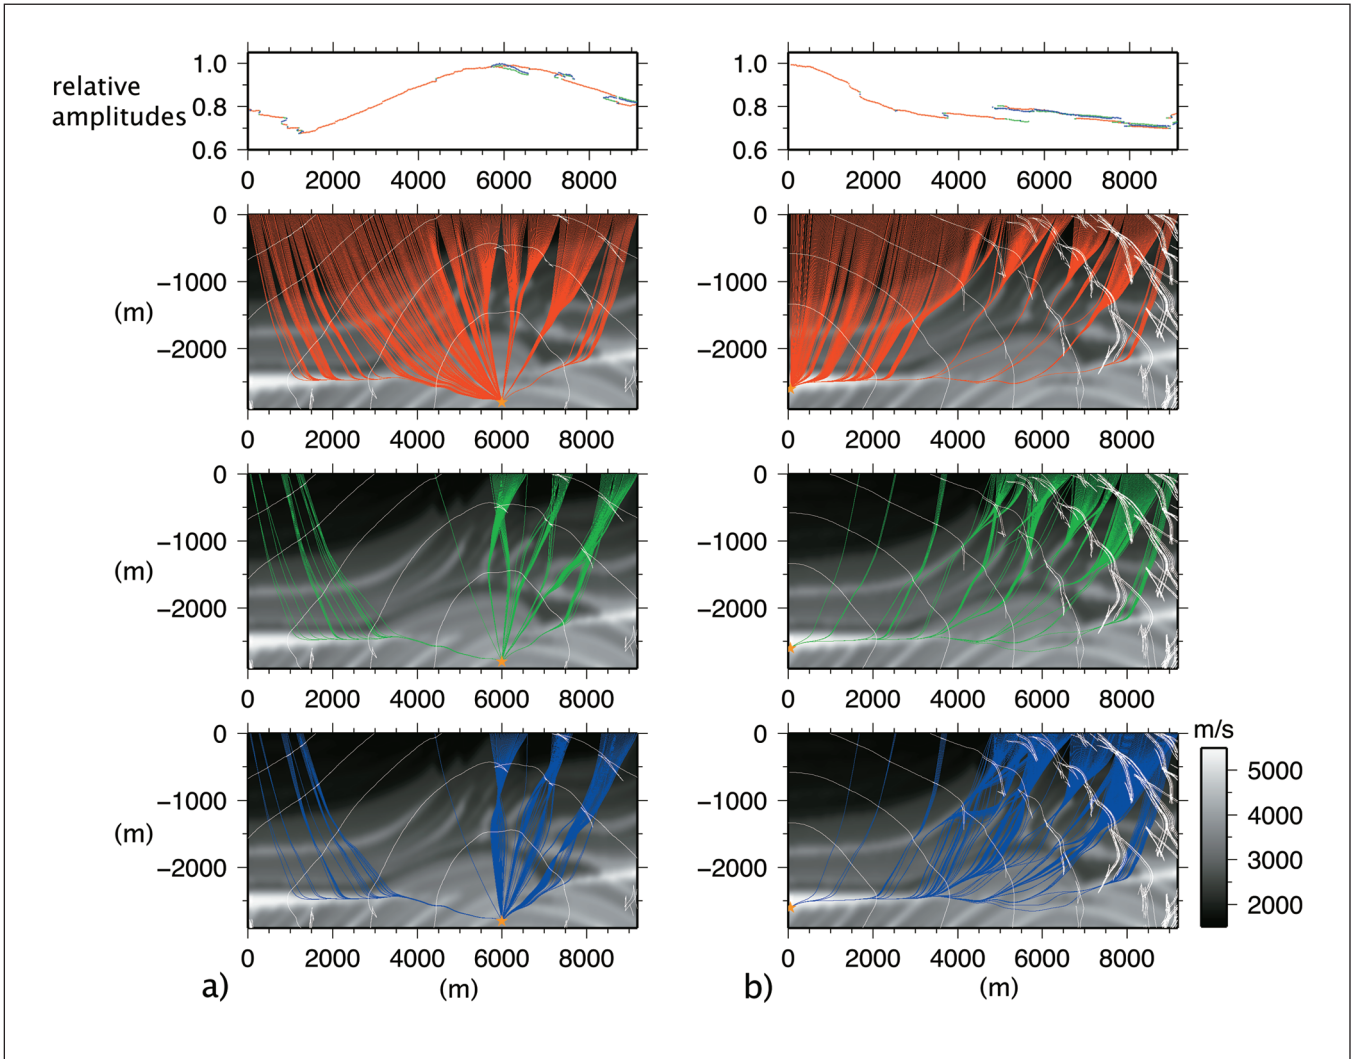
Fig. 7. Later arrivals and amplitude information for two sources (a and b) in the Marmousi model. The raypaths for the first arrivals are given in red, for the second arrivals in green, and for third and later arrivals in blue. The figure on the top shows the relative amplitudes for the first (red), second (green), and third arrival (blue).

In the Eulerian scheme, the signed distance functions are defined on a grid consisting of 113×58×59 nodes. The grid is scaled so that the spacing in all three dimensions is similar. Calculating the wavefronts using the Eulerian solver takes 24 h compared with 14 s when using the Lagrangian solver. In the Lagrangian solver we represent the point source using 500 points. As shown in Figures 6(b) and 6(c), the wavefronts calculated by the Lagrangian solver are clearly preferable. Without a finer computational mesh the level set method is unable to capture the development of the swallowtail pattern. Increasing the grid size by a factor of two, however, means that the computation time increases by a factor of sixteen and the memory requirements increase by a factor of eight. The Lagrangian solver predicts three arrivals for the receiver and we also can extract the ray-path information for each arrival.

It is important to realise that we are computing later arrivals using a velocity model that was constructed using only first arrivals. This may mean that later arrivals are under-represented, as first-arrival tomography tends to underestimate the amplitude of slow anomalies. Figure 6(c) suggests that the inclusion of later arrivals in seismic imaging (e.g. in tomography) could in principle improve path coverage and hence resolution, since raypaths for the two later arrivals are sampling different parts of the model than the first arrival.

Later Arrivals in the Marmousi Model

The Marmousi model has traditionally been used to demonstrate the limitations of first-arrival traveltimes in imaging complex media. Geoltrain and Brac (1993) showed that multi-arrival traveltimes are needed in order to accurately image the Marmousi model. The geological structure is based on a profile through the North Quenguela trough in the Cuanza basin in Angola (Versteeg, 1993).

Traditionally traveltimes have been computed using a smoothed version of the Marmousi model (see <http://www.rocq.inria.fr/~benamou/testproblem.html>) (e.g., Buske and Kaestner, 2004; Qian and Leung, 2004; Coman and Gajewski 2005). The smooth model was obtained by convolving the so-called hard model, which is characterised by strong velocity gradients, with a spatial Hanning (cos²) filter of radius 150 m. The model consists of 384×122 nodes with a vertical separation of 24 m. Figure 7(a) illustrates raypaths, wavefronts, and relative amplitudes obtained by the Lagrangian solver for the standard test problem, using the smoothed version of the model with a source at (6000, -2800). We are able to recover 651 arrivals and the corresponding raypaths for the 384 receivers positioned at the surface. The computation time in this case is 45 s, and more than 40% of the arrivals are later arrivals, which tend to sample regions avoided by first arrivals. Figure 7(b) shows the raypaths for a source close to the

left boundary of the model, at (50, -2600). As the wavefronts for this source position travel a greater distance, the wavefronts become much more complex and the fast region acts as a kind of waveguide. We observe 3291 arrivals for this example, with later arrivals accounting for more than 85% of the total. For both source locations we observe that the later arrivals tend to sample regions avoided by the first arrivals. We also calculate relative amplitudes for the first, second, and third arrivals using dynamic ray tracing (e.g., Vinje et al., 1993). The later arrivals have similar, and for some receivers, larger amplitudes than the first arrival. This confirms our expectation that the first arrival of a wavefront is not always the most energetic arrival.

We use the Lagrangian solver to identify regions in the Marmousi model where sources generate later arrivals for at least one receiver on the surface. We do this by evaluating 2976 source points with a vertical and horizontal spacing of 48 m. Figure 8(a) shows a map of the maximum difference in arrival time between the first and second arrival at one of the receivers at the surface. Sources close to the receivers in the centre of the model show a large difference in traveltime between the first and second arrival. This is due to the strong velocity gradient in the lower left region of the model. This strong change in the velocity behaves like a reflector for sources above this fast region and deflects down-going energy back towards the surface. Figure 8(b) is a map of the maximum number of later arrivals at the surface, as a function of source location. Note that for reasons of computational convenience we limit the maximum number of arrivals to 60. Sources close to the left and right boundary of the model tend to generate large numbers of later arrivals, as the wavefronts can travel greater distances and therefore more, and larger, swallowtail patterns can develop. If a source is located in the two fast regions near to the left and right edges of the model, we do not observe large numbers of later arrivals at the surface. A wavefront in the neighbourhood of a point source tends not to develop a swallowtail, even in the presence of a significant velocity contrast, as a result of its high curvature.

Figure 8(c) is a map of the velocity gradient of the Marmousi model. By comparing the wavefronts in Figure 7 with Figure 8(c) we see that the swallowtails are initiated in regions where we observe a strong velocity gradient. Overall we observe that the generation and detection of multiple arrivals is extremely sensitive to small changes in the velocity structure and the source and receiver location.

DISCUSSION AND CONCLUSIONS

We have given an overview of two classes of phase space solvers for the computation of multiple arrivals: a Lagrangian and an Eulerian scheme. The fundamental difference between the two solvers is their scaling with computation time, memory, and accuracy. The Lagrangian approach uses a local one-dimensional discretisation, while the Eulerian scheme requires a three-dimensional discretisation even for a two-dimensional problem.

The level set method is well known and frequently used in computational mathematics, but little known or used in seismology. Even so there is still a lack of agreement in the literature on how best to implement it. The stability of solvers used for the level set equation, and the implementation of schemes aimed at keeping the signed distance function well-behaved during the computation, form part of ongoing research. The level set method has several powerful features for general interface problems, among them the capability to handle the merging and breaking of interfaces automatically. However the fundamental limitation of the level set method is that features of a surface that are smaller than a few grid cells cannot be adequately resolved by a signed distance function. In practice
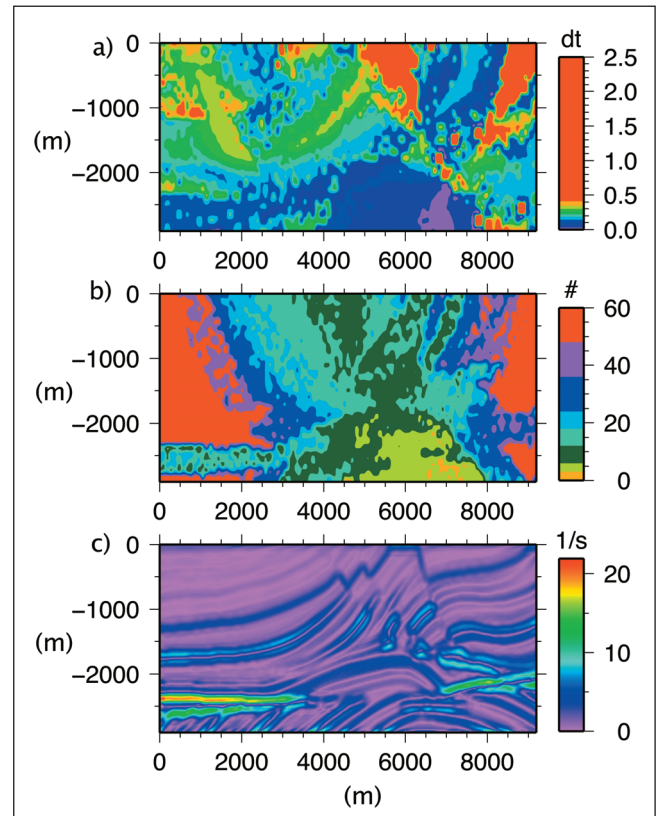

Fig. 8. (a) Maximum observed traveltime difference between first and second arrival for any one of the receivers at the surface as a function of source location. (b) Maximum number of arrivals for any one of the receivers at the surface as a function of source location. (c) Velocity gradient of the Marmousi model.

the grid resolution in three dimensions for the representation of the bicharacteristic strip will therefore always be a limiting factor. A narrow-band adaptive gridding technique for the signed distance function could provide a high grid resolution where it is necessary. However such algorithms are generally difficult to implement and it is questionable if such a scheme could ever be superior to a Lagrangian scheme for seismological problems.

Lambaré et al. (1996) use the two components of the slowness vector as the additional coordinates for their phase space, while we only use the direction of the local wavefront normal as an additional coordinate. This means that we have a reduced three-dimensional phase space when compared to their four-dimensional phase space. Coman and Gajewski (2005) do not propagate their wavefront in phase space and therefore calculate an approximation to its curvature based on the local slowness vector in order to produce a satisfactory refinement criterion. Our approach also differs from previous implementations in the sense that we remove points from the wavefront if their density is above a certain threshold. We are able to track more than 60 later arrivals for certain source locations in the Marmousi model because of the stability of the method. Coman and Gajewski (2005) and Lambaré et al. (1996) also calculate later arrivals for this model but because of the position of their sources they track fewer later arrivals in their examples.

A phase-space-based approach for the computation of later arrivals in a three-dimensional model would require the evolution of a surface in five dimensions. The two phase-space coordinates are then given by the two angles describing the direction of the local wavefront normal. In the Eulerian approach the level set method would have to be used to evolve three surfaces in five dimensions (Osher et al., 2002). In this case the grid resolution

in the Eulerian approach would be severely limited by the computational cost. For the Lagrangian case, extension to three dimensions would mean representing a two-dimensional manifold in a five dimensional phase space, by triangles for example, and keeping track of the connectivity between these triangles, which seems much more feasible. Vinje et al. (1996) performed wavefront tracking in three dimensions, but the refinement of the wavefront is not performed in phase space, and they do not remove triangles when the wavefront becomes over sampled.

In summary, we have shown that it is possible to compute multi-arrival traveltimes in a complex two-dimensional velocity model and that the Lagrangian scheme is currently the preferred choice by a wide margin. We have also demonstrated that the practicality of the level set method for computing multi-arrival traveltimes is currently questionable, despite the level set method being a well established and proven technique outside of seismology. We used the Lagrangian solver in order to find source locations for the Marmousi model that will lead to the observation of large numbers of later arrivals at a given set of receivers. We have also shown how the Lagrangian solver can be used to find receiver positions that will detect later arrivals. A longer-term goal is to exploit the information contained in observed later arrivals in order to improve seismic imaging.

REFERENCES

Benamou, J.D., 2003, An introduction to Eulerian geometrical optics (1992–2002): *Journal of Scientific Computing*, **19**, 63–93.

Benamou, J.D., 1999, Direct computation of multivalued phase space solutions for Hamilton-Jacobi equations: *Communications on Pure and Applied Mathematics*, **52**, 1443–1475.

Buske, S., and Kaestner, U., 2004, Efficient and accurate computation of seismic travel times and amplitudes: *Geophysical Prospecting*, **52**, 313–322.

Cerveny, V., 2001, *Seismic Ray Theory*: Cambridge University Press, 1st edition.

Chapman, C.H., 1985, Ray theory and its extensions - WKBJ and Maslov seismograms: *Journal of Geophysics*, **58**, 27–43.

Cockburn, B., Qian, J., Reitich, F., and Wang, J., 2005, An accurate spectral/discontinuous finite-element formulation of a phase-space-based level set approach to geometrical optics: *Journal of Computational Physics*, **208**, 175–195.

Coman, R., and Gajewski, D., 2005, Traveltime computation by wavefront-orientated ray tracing: *Geophysical Prospecting*, **53**, 23–26.

Conder, J.A., and Wiens, D.A., 2006, Seismic structure beneath the Tonga arc and Lau backarc basin determined from joint Vp, Vp/Vs tomography: *Geochemistry Geophysics Geosystems*, **7**, Art. Q03018.

Fatemi, E., Engquist, B., and Osher, S., 1995, Numerical solution of the high frequency asymptotic expansion for the scalar wave equation: *Journal of Computational Physics*, **120**, 145–155.

Fomel, S., and Sethian, J.A., 2002, Fast-phase space computation of multiple arrivals: *Proceedings of the National Academy of Sciences of the United States of America*, **99**, 7329–7334.

Geoltrain, S., and Brac, J., 1993, Can we image complex structures with first arrival travel time?: *Geophysics*, **58**, 564–575.

Harten, A., Engquist, B., Osher, S., and Chakravarthy, S.R., 1987, Uniformly high order accurate essentially non-oscillatory schemes, iii: *Journal of Computational Physics*, **71**, 231–303.

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

Jiang, G.S., and Peng, D.P., 2000, Weighted ENO schemes for Hamilton-Jacobi equations: *SIAM Journal on Scientific Computing*, **21**, 2126–2143.

Jiang, G.S., and Shu, C.W., 1996, Efficient implementation of weighted ENO schemes: *Journal of Computational Physics*, **126**, 202–228.

Kim, S., and Cook, R., 1999, 3-D traveltime computation using second-order ENO scheme: *Geophysics*, **64**, 1867–1876.

Lambaré, G., Lucio, P.S., and Hanyga, A., 1996, Two-dimensional multivalued traveltime and amplitude maps by uniform sampling of a ray field: *Geophysical Journal International*, **125**, 584–598.

Lambaré, G., Virieux, J., Madariaga, R., and Jin, S., 1992, Iterative asymptotic inversion in the accoustic approximation: *Geophysics*, **57**, 1138–1154.

Liu, X., Osher, S., and Chan, T., 1994, Weighted essentially non-oscillatory schemes: *Journal of Computational Physics*, **115**, 200–212.

Lucio, P.S., Lambaré, G., and Hanyga, A., 1996, 3D multivalued travel time and amplitude maps: *Pure and Applied Geophysics*, **148**, 449–479.

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

Osher, S., and Fedkiw, R., 2003, *Level Set Methods and Dynamic Implicit Surfaces*: Springer Verlag, 2nd edition.

Osher, S., Cheng, L.T., Kang, M., Shim, Y., and Tsai, Y.H., 2002, Geometric optics in a phase-space-based level set and Eulerian framework: *Journal of Computational Physics*, **179**, 622–648.

Osher, S., and Sethian, J., 1988, Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations: *Journal of Computational Physics*, **79**, 12–49.

Pereyra, V., Lee, W.H.K., and Keller, H.B., 1980, Solving two-point seismic-ray tracing problems in a heterogeneous medium. Part 1. a general adaptive finite difference method: *Bulletin of the Seismological Society of America*, **70**, 79–99.

Qian, J., and Leung, S., 2004, A level set based Eulerian method for paraxial multivalued traveltimes: *Journal of Computational Physics*, **197**, 711–736.

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

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

Sethian, J.A., 1999, *Level Set Methods and Fast Marching Methods*: Cambridge University Press, 2nd edition.

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

Sheriff, R.E., and Geldart, L.P., 1995, *Exploration Seismology*: Cambridge University Press, 2nd edition.

Shu, C., and Osher, S., 1988, Efficent implementation of essentially non-oscillatory shock capturing schemes: *Journal of Computational Physics*, **77**, 439–471.

Shu, C., and Osher, S., 1989, Efficent implementation of essentially non-oscillatory shock-capturing schemes, ii: *Journal of Computational Physics*, **83**, 32–78.

Sun, Y., 1992, Computation of 2D multiple arrival traveltime fields by an interpolative shooting method: *SEG Annual Meeting Technical Program Expanded Abstracts*, **62**, 1320–1323.

Veersteg, J.R., 1993, Sensitivity of prestack depth migration to the velocity model: *Geophysics*, **56**, 873–882.

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

Vidale, J.E., 1990, Finite-difference calculation of traveltimes in 3 dimensions: *Geophysics*, **55**, 521–526.

Vinje, V., Iversen, E., Åstebøl, K., and Gjøystdal, H., 1996, Estimation of multivalued arrivals in 3D models using wavefront construction - Part I: *Geophysical Prospecting*, **44**, 819-842.

Vinje, V., Iversen, E., and Gjoystdal, H., 1993, Travel-time and amplitude estimation using wave-front construction: *Geophysics*, **58**, 1157–1166.

Zelt, C.A., and Ellis, R.M., 1988, Practical and efficent ray tracing in two-dimensional media for rapid traveltime and amplitude forward modelling: *Journal of the Canadian Society of Exploration Geophysicists*, **24**, 16–31.