*Geophys. J. Int.* (1991) **105**, 271–284

# Finite difference computation of traveltimes in very contrasted velocity models: a massively parallel approach and its associated tools

Pascal Podvin[1] and Isabelle Lecomte[2]

[1] *CGGM/Géophysique, Ecole des Mines de Paris, Fontainebleau, France*
[2] *Laboratoire de Sismologie, Institut de Physique du Globe, Strasbourg, France*

Accepted 1990 November 14. Received 1990 November 13; in original form 1990 September 10



## SUMMARY

We present a new massively parallel method for computation of first arrival times in arbitrary velocity models. An implementation on conventional sequential computers is also proposed.

This method relies on a systematic application of Huygens' principle in the finite difference approximation. Such an approach explicitly takes into account the existence of different propagation modes (transmitted and diffracted body waves, head waves). Local discontinuities of the time gradient in the first arrival time field (e.g., caustics) are built as intersections of locally independent wavefronts. As a consequence, the proposed method provides accurate first traveltimes in the presence of extremely severe, arbitrarily shaped velocity contrasts.

Associated with a simple procedure which accurately traces rays in the obtained time field, this method provides a very fast tool for a large spectrum of seismic and seismological problems.

We show moreover that this method may also be used to obtain several arrivals at a given receiver, when the model contains reflectors. This possibility significantly extends the domain of potential geophysical applications.

**Key words:** finite difference, massively parallel computation, ray tracing, seismic waves, traveltimes.

## INTRODUCTION

Traveltimes for seismic waves are classically computed with ray tracing techniques. Conversely, graphical methods devoted to wavefront tracing on simple examples were proposed by Thornburgh as early as 1930, and further generalized by Riznichenko (1946) to the case of layered media. Recently, Vidale (1988, 1990) proposed a general wavefront tracing technique based on a finite difference approximation of the eikonal equation. Quickly acknowledged as a new and promising approach, Vidale's first paper opened the way to several recent studies addressing the same question in a variety of ways. Saito (1989, 1990) and Moser (1989) use the graph theory, whereas Van Trier & Symes (1990, 1991) solve a Hamiltonian representation of the eikonal equation. Qin *et al.* (1990) propose an improvement of Vidale's original algorithm.

These methods are only devoted to computation of first traveltimes. According to Van Trier & Symes (1991), this limitation intrinsically comes from mathematical properties of the eikonal equation. As a consequence, on a given model, the finite difference approach should fail to provide several arrival times at a given receiver. Although this statement relies on an unproved mathematical conjecture, it implies that Vidale's approach should certainly not be considered as a mere substitute to ray tracing. Notwithstanding this limitation, a large spectrum of potential applications of a fast first traveltime computation method to seismic and seismological problems was immediately identified by Vidale. In practice, first applications published were devoted to earthquake location (Nelson & Vidale 1990).

However, all existing finite difference methods up to now encounter serious difficulties when applied to models containing sharp first-order velocity contrasts. Moreover, as mentioned by Van Trier & Symes, Vidale's algorithm cannot be implemented on vectorial or parallel computers, although some applications of this method (e.g., to seismic tomography) might be computationally very intensive.

In this paper, we propose a new finite difference

algorithm which overcomes both difficulties. This method may be applied to virtually any velocity model in 2-D as well as 3-D. Step velocity contrasts as high as 1:10, whatever their shape, are correctly taken into account. Moreover, this algorithm is explicitly designed as a massively parallel scheme, although it may also simply be implemented on a conventional computer.

In the following section, we introduce the computational method, and discuss its accuracy. *A posteriori* ray tracing is then described. Some associated tools designed to extend the possible uses of such method are then discussed. Finally, we briefly present a few examples of geophysical applications.

## METHOD

This section presents a new finite difference algorithm devoted to computation of first arrival times in virtually any velocity model. Although this method is inspired from Vidale's (1988, 1990) approach, its basic principles are fundamentally different.

### Physical principles for first arrival times computation

In the ray approximation, wave propagation is described by the eikonal equation:

$$(\nabla t)^2 = s^2$$

where $s(\mathbf{x})$ is the slowness of the medium (inverse of wave velocity) and $t(\mathbf{x})$ represents the arrival time of a wavefront at point $\mathbf{x}$. According to the classical ray approximation (asymptotic approximation of the wave equation at infinite frequency), wavefronts are treated as propagating discontinuities (represented as isochrons). Vidale's approach relies on a finite difference approximation of the eikonal equation. With this approach, wavefronts rather than rays are propagated in the model. In the following paragraphs, we shall use alternately the ray viewpoint and the wavefront viewpoint, although, basically, the method does not involve the notion of ray. Both viewpoints carry exactly the same information, i.e., a description of arrival time field $t(\mathbf{x})$.

Computation of first arrival times with ray tracing in heterogeneous media is difficult in part because any point in the model receives an unpredictable number of rays, among which only one is to be picked. In many cases, rays traced from the source may never reach certain regions of the medium (shadow zones), where signals received are diffracted by some localized secondary source. Ray tracing in shadow zones may then only be achieved at the expense of an explicit diagnosis of diffraction. The use of finite differences does not suppress these problems: a change in mathematical method does not change the nature of physical difficulties to be solved, but may only provide more efficient tools.

This simple preliminary remark explains almost totally why Vidale's algorithm fails in very contrasted models. Without further precautions, a plain finite difference approximation of the eikonal equation implicitly consists in mathematically propagating a single wavefront in the model, whereas several locally independent wavefronts may in fact reach any point. As a consequence, the algorithm silently mixes information related to these wavefronts, which may lead to mathematically conflicting situations (in practice, negative square root arguments). For this reason, the proposed algorithm is not, as Vidale's method, a simple finite difference approximation of the eikonal equation in the sense that multiple arrivals at any point are systematically considered. A first arrival criterion is then used in order to pick the first one.

Such approach explicitly takes into account the fact that first arrival time fields usually contain discontinuities of the time gradient (e.g., caustics) so that they may not be seen, from a mathematical viewpoint, as regular solutions of the eikonal equation. The solution adopted here does not involve the mathematical conjecture proposed by Van Trier & Symes (1991) according to which the generalized (viscosity) solution of the eikonal equation picked by plain finite differences is the first arrival time field. With our approach, any discontinuity of the time gradient explicitly results from the intersection of several regular solutions (i.e., several independent wavefronts).

### Physical representation of the slowness model

In order to understand properly the physics of propagation in the model, it is necessary to adopt a physical representation of the discrete model on which computations take place. In opposition with Vidale's description, we shall consider the discretization of the slowness model as a physical approximation: the real model is replaced by a virtual model in which space is paved by constant velocity materials. Computation takes place on this approximate model, so that problems of propagation are non-trivial only at mesh boundaries.

The accuracy of this initial approximation essentially depends on mesh spacing. Because, in practice, this method will be applied to band-limited signals, mesh spacing should be chosen according to the actual wavelength of the real signal, which, in turn, constrains the accuracy of ray approximation in a given model, or, conversely, the acceptable spatial resolution of model description. As an initial rule, mesh spacing should be at least one order of magnitude smaller than signal wavelength which, in turn, should be smaller than characteristic dimensions of slowness anomalies considered.

On the basis of this physical representation, the problem to be solved may now be formulated unambiguously: consider a model consisting of cubic meshes (or square meshes in 2-D) in which the slowness $s$ is constant. Let $h$ be the mesh spacing. A source located anywhere in the model emits an impulse at time $t = 0$. At what time does each grid-point receive the first signal? This formulation implies that for $n_x n_y n_z$ meshes, $(n_x + 1)(n_y + 1)(n_z + 1)$ times are to be computed.

Two steps are required in order to build the corresponding computational scheme. The use of finite differences implies that computations are to be thought of as local: the arrival time at a given grid-point only depends on arrival times at its neighbours and the local values of slowness. In the first step, this local computation must be designed. The second step will address the order in which arrival times are computed, i.e., how do we propagate computations? Vidale only defined conventional marching

procedures: we shall show that different strategies are available.

### Local computation of arrival times

*Huygens' principle and plane wavefront approximation*

Figure 1 shows the data structures attached to a point in 2-D (eight neighbours, four values of slowness) and 3-D (26 neighbours, eight values of slowness). The local problem may be formulated as follows: considering that first arrival times are known at every neighbour, how do we estimate first arrival time at the current grid-point by finite differences? This question is to be solved on a purely local basis, without any reference to global wavefront propagation. In particular, the position of the source with respect to the current point needs not be known, even approximately.

This local problem may be solved by the means of Huygens' principle. As shown in Fig. 1, we shall consider that the current point is surrounded by a continuous set of secondary sources. The loci of these sources define a square in 2-D, a cubic box in 3-D. Huygens' principle states that every point on these boundaries acts as a secondary source emitting an impulse at the moment it is reached by the first arrival. The principle of local computation thus consists of timing all arrivals at the current point from this local set of Huygens' sources, and picking the first one.

For this purpose, one more approximation is needed because only very few of these secondary sources are actually timed, i.e., the current point's neighbours, whereas all of them should be properly timed in order to apply Huygens' principle. In 2-D, we timed sources between two neighbour grid-points by linear interpolation between their allocated times. In 3-D, we used the same approach, but then, several interpolations are possible because interpolation is overdetermined by the knowledge of arrival times at the four corners of each interface. This overdetermination could be avoided by the choice of a more complex interpolation (e.g., polynomial), but this choice would imply some sort of smoothing of the local time field. As none of these interpolations appears physically better, we adopted the most simple (and computationally economic) solution: at any point on an interface, several linear interpolations between times at three corners were estimated and the lowest one was picked. Another advantage of this solution is its very simple interpretation in terms of a physical approximation. Linear interpolation between grid-points may be seen as a local plane wavefront approximation. Fig. 2 shows in what sense these approximations are linked.

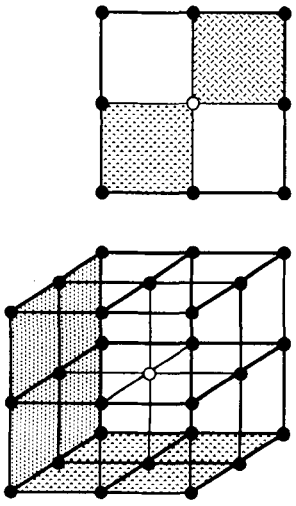

**Figure 1.** Data structures used in 2-D and 3-D. First arrival time is computed at the central grid-point (current point). Boundaries of the surrounding box are considered as a continuous set of Huygens' sources.

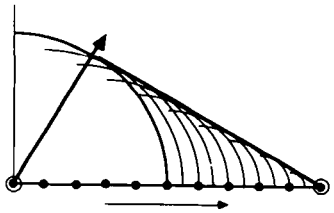

**Figure 2.** Sketch of wavelets emitted by an interface in 2-D. The corresponding wavefront is built as their envelope. When intermediate secondary sources are timed by linear interpolation, this wavefront is partly plane ('illuminated' zone), partly cylindrical (single wavelet emitted by an extremity of the interface).

For this reason, and for a matter of clarity, we shall now abandon Huygens' principle as an explanating tool of the local computational scheme and rather describe it in terms of propagation of locally plane wavefronts. These two explanations are equivalent, but the following one appeals to more classical views in terms of propagation in the ray approximation. In particular, it will illustrate how this approach explicitly deals with different propagation modes, i.e., body waves, head waves, diffraction by point diffractors.

Clearly, the local plane wavefront approximation fails at the vicinity of the source, but this imprecision will only be discussed later, when initialization procedures are considered. In the following paragraphs, the source is implicitly considered to be far from the current point.

### Transmission stencils in 1-D, 2-D and 3-D

In the discrete model, transmission through interfaces is by far the most frequent process by which energy travels from a mesh to its neighbours. For the sake of clarity, we shall first examine the situation in 2-D.

*2-D transmission.* In 2-D, the time difference between two neighbour grid-points M and N refers to the transmission of a locally plane wavefront through interface MN. As shown in Fig. 3, the knowledge of this time difference leaves an ambiguity on the direction of propagation. If $s$ is the value of slowness in one of the adjacent meshes, two estimates of the time gradient satisfying the eikonal equation in this mesh may be computed:

$$\frac{\partial t}{\partial x} = \frac{t_N - t_M}{h}, \qquad \frac{\partial t}{\partial y} = \pm \sqrt{s^2 - \frac{(t_N - t_M)^2}{h^2}}.$$

The missing information is the sign of $\partial t/\partial y$, which entirely defines whether the wavefront is transmitted to the current mesh, or on the contrary, to its neighbour. When this information is known, an estimate of arrival time at

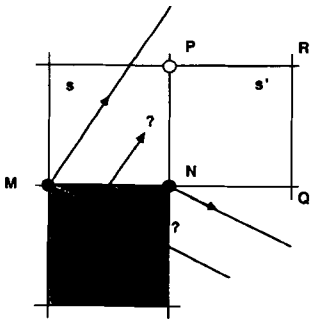


**Figure 3.** Transmission in 2-D. When only points M and N are timed, the direction of propagation is not entirely defined. Moreover, if point P is illuminated by the transmitted plane wavefront, slowness $s'$ in mesh NPQR may be so low as to imply the generation of a head wave from N to P.

point P (see Fig. 3) may be computed if point P is reached by this wavefront, i.e. if

(i) $0 \leq (t_N - t_M) \leq hs/\sqrt{2}$,
(ii) $\partial t/\partial y \geq 0$, then

$$t_P = t_N + \sqrt{(hs)^2 - (t_N - t_M)^2}.$$

This fundamental 2-D transmission stencil is conditional: an 'illumination condition' must be verified in order to assess transmission from the interface to the current point. This condition tests whether (i) a ray connects interface MN to point P, and (ii) actual propagation (i.e., slowness vector) is oriented from MN to P. When put in terms of Huygens' principle, the only condition which appears strictly necessary is the first one. This is a very important simplification, which may also be interpreted as a consequence of Fermat's principle. An error on the sign of $\partial t/\partial y$ is harmless because, in such case, the estimated arrival time in P results from integration of slowness along a path which violates the eikonal equation. Fermat's principle then states that this arrival time is certainly superior or equal to the first arrival time at point P. For this reason, and under the condition that all relevant estimates of arrival times at point P are exhaustively computed, the condition on the sign of $\partial t/\partial y$ may be abandoned. As a consequence, the transmission stencil in 2-D is entirely defined by the knowledge of a single time difference.

In the 2-D data structure defined in Fig. 1, eight different estimates of arrival time at the current point are potentially to be computed on the basis of 2-D transmission.

*1-D transmission: generation and propagation of head waves in 2-D.* Suppose, as in Fig. 3, that arrival time $t$ at point P was estimated by 2-D transmission through interface MN. Clearly, $t$ does not depend on the value $s'$ of slowness in the adjacent mesh NPQR. In particular, $s'$ may be such that:

$$t_P - t_N > hs'.$$

Such a relation simply means that the incidence of the transmitted wavefront on interface NP exceeds the critical angle. In this case, a head wave travelling along interface NP with slowness $s'$ must be generated (or transmitted). As a consequence, point P is reached by an earlier arrival at time

$$t_P = t_N + hs'.$$

This relation defines a 1-D transmission stencil which physically deals with generation and transmission of interface waves in 2-D. A precaution must thus be added in order to prevent artifacts due to spurious head wave generation at the frontiers of the model: slowness is treated as infinite at the model's boundaries (analogous to an absorbing boundary condition). With this precaution, the 1-D transmission stencil is defined as follows: suppose that arrival time is known at point N, and let P be a first neighbour of N, $s$ and $s'$ the values of slowness in the two adjacent meshes, then the 1-D transmission estimate of arrival time at point P is

$$t_P = t_N + h \min(s, s').$$

This stencil is unconditional because, as explained earlier, Fermat's principle states that if no head wave actually propagates from N to P, then this time estimate is certainly superior to the first arrival time we are seeking.

In the 2-D data structure, four independent 1-D estimates must thus be computed as potential first arrival times at the current point.

*Transmission in 3-D.* In 3-D, in order to stick to the plane wavefront approximation, several independent plane wavefronts are considered to be transmitted through each interface. Fig. 4 shows that four stencils based on three different topologies are applicable for estimation of arrival time at point R by 3-D transmission through interface MNPQ. Each stencil implicitly deals with the transmission of a locally plane wavefront through one half of the interface. Each plane wavefront is defined by arrival times at three corners of the transmitting interface. As in 2-D, these stencils are conditional, the condition being that a ray path actually links the current point R to the half interface considered. Thanks to Fermat's principle, the sign of $\partial t/\partial z$ is not tested so that each stencil only depends on three arrival times.

The four stencils are:

MNP → R: if

$$t_M \leq t_N, \quad t_M \leq t_P,$$
$$2(t_P - t_M)^2 + (t_N - t_M)^2 \leq (hs)^2,$$
$$2(t_N - t_M)^2 + (t_P - t_M)^2 \leq (hs)^2,$$
$$(t_N - t_M)^2 + (t_P - t_M)^2 + (t_N - t_M)(t_P - t_M) \geq (hs)^2/2,$$

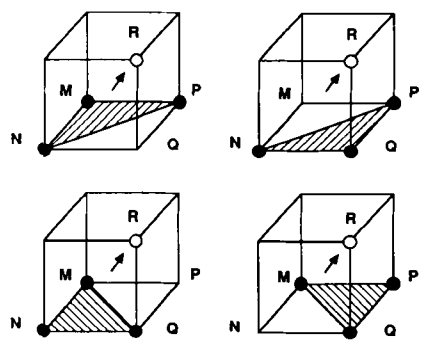


**Figure 4.** Transmission in 3-D. Four locally plane wavefronts are considered to be transmitted through interface MNPQ. Each wavefront involves arrival times at three corners of the interface.

then

$$t_R = t_N + t_P - t_M + \sqrt{(hs)^2 - (t_N - t_M)^2 - (t_P - t_M)^2};$$

QNP→R: if

$$t_N \leq t_Q, \quad t_P \leq t_Q,$$

$$(t_Q - t_N)^2 + (t_Q - t_P)^2 + (t_Q - t_N)(t_Q - t_P) \leq (hs)^2/2,$$

then

$$t_R = t_Q + \sqrt{(hs)^2 - (t_Q - t_N)^2 - (t_Q - t_P)^2};$$

NMQ→R: if

$$0 \leq (t_N - t_M) \leq (t_Q - t_N),$$

$$2(t_Q - t_N)^2 + (t_N - t_M)^2 \leq (hs)^2,$$

then

$$t_R = t_Q + \sqrt{(hs)^2 - (t_Q - t_N)^2 - (t_N - t_M)^2};$$

PMQ→R: if

$$0 \leq (t_N - t_M) \leq (t_Q - t_N),$$

$$2(t_Q - t_N)^2 + (t_N - t_M)^2 \leq (hs)^2,$$

then

$$t_R = t_Q + \sqrt{(hs)^2 - (t_Q - t_P)^2 - (t_P - t_M)^2}.$$

In the 3-D data structure, 24 interfaces are considered, so that up to 96 independent transmission stencils provide potential first arrival times at the current point.

*Head wave generation and propagation in 3-D.* Head wave generation and propagation in 3-D is practically identical to the 2-D problem embedded in 3-D. As in 2-D, any potential arrival at the arrival may be accurately picked. As a consequence, 2-D and 1-D transmission stencils defined earlier are also applicable to the 3-D case. According to their physical interpretation, the 24 2-D transmission stencils in 3-D use the lower value of slowness in the two adjacent meshes. In the same way, the six 1-D transmission stencils use the lowest value of slowness in the four meshes concerned.

### *Diffracted wavefronts*

Figure 5 shows a 2-D wavefront pattern in which point P is located in a local shadow zone. Two wavefronts are actually transmitted to the current mesh, but none of them reaches point P. In the sense of finite differences, the physical interpretation of this pattern is immediate: point M as a point diffractor. P receives a diffracted wavefront emitted by this secondary source, so that the corresponding arrival time is

$$t_P = t_M + hs\sqrt{2}.$$

This relation defines the 2-D diffraction stencil. Once more, it can be computed systematically because its result is superior or equal to the first arrival time at the current point. This 2-D stencil must also be used in the 3-D case, in order to complete all aspects of wave propagation along interfaces (head waves). In this case (Fig. 6), the lower value of slowness in the two meshes adjacent to the interface considered is used.

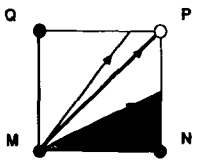


**Figure 5.** Local shadow zone in 2-D. Corner M acts as a secondary source.

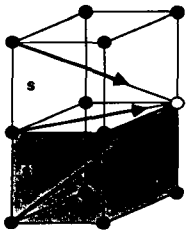


**Figure 6.** Diffracted body waves and diffracted head wave in 3-D.

The 3-D diffraction stencil (Fig. 6) consists in direct arrival to the current point from any of the eight corners of the 3-D data structure. It must be systematically computed in order to complete the examination of all possible arrivals at the current point. With the notations defined in Fig. 4:

$$t_R = t_M + hs\sqrt{3}.$$

Edge diffraction must also be taken into account in 3-D. 24 diffracting edges are thus considered in the 3-D data structure. As an example, edge MN in Fig. 4 must be considered as a set of diffracting secondary sources when point R is timed. The corresponding stencil picks the earliest arrival from this set. The associated condition simply ensures that the picked arrival comes from a point located within edge limits: if

$$0 \leq (t_N - t_M) \leq hs/\sqrt{3}$$

then

$$t_R = t_N + \sqrt{2}\sqrt{(hs)^2 - (t_N - t_M)^2}.$$

### *Summary: local computation in 2-D and 3-D*

The local computational scheme may now be summarized as follows.

In 2-D, when arrival times are known at each of the eight neighbours, 16 independent stencils are used: four 1-D transmitted arrivals (head waves), eight 2-D transmitted plane wavefronts (conditional), and four 2-D diffracted arrivals.

In 3-D, on the basis of 26 known arrival times at current point's neighbours, up to 170 stencils are applicable: six 1-D transmitted arrivals, 24 2-D transmitted arrivals (conditional), 12 2-D diffracted arrivals (1-D and 2-D stencils deal with head waves), 96 3-D transmitted plane wavefronts (conditional), and 32 3-D diffracted arrivals.

The local computation simply consists of computing all these time estimates and returning the lowest one. Apart from 2-D and 3-D transmission stencils, all other time estimates may be seen as direct arrivals from each neighbour. Their computation is immediate so that computational cost is mostly due to transmission stencils (square root extraction). In practice, however, most illumination conditions are usually not verified, so that only a few transmitted estimates are actually computed at each point.

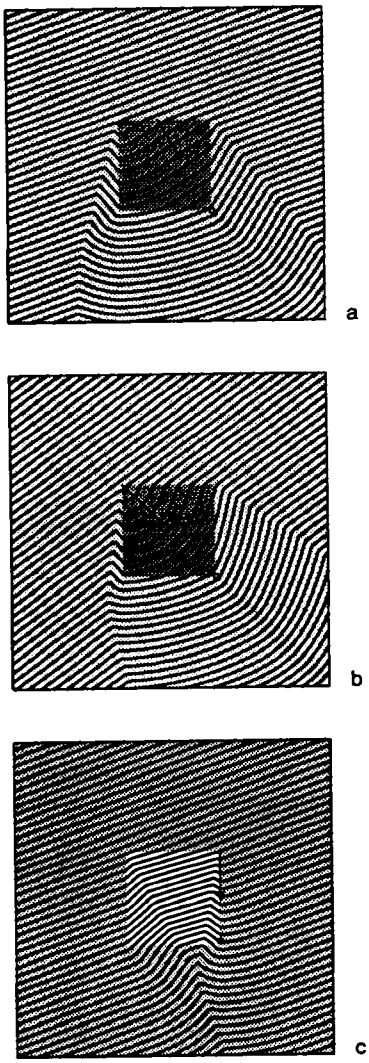


**Figure 7.** Wavefront patterns in 2-D. A plane wavefront propagating from the upper left region of the model is incident on a square slowness anomaly. Three examples illustrate how arrival time is evaluated at the lower right corner. When the anomaly is fast (a and b), 2-D transmission or diffraction stencils are involved according to angle of incidence. In (c), the anomaly is slow, so that a head wave is generated along the right side of the anomaly.

Figure 7 shows a few typical 2-D patterns which are commonly encountered in a mesh during local computation. This figure summarizes the role and physical interpretation of the different stencils used. In Figs 7(a) and (b), a plane wavefront is incident on a fast anomaly. According to the angle of incidence, the first arrival time at the lower right corner of the anomaly results from transmission through an interface (Fig. 7a) or diffraction by a corner (Fig. 7b). In Fig. 7(c), the anomaly is slow so that the first arrival time at the same point results from a head wave travelling along an interface with the velocity of the fast medium.

## Propagation of computations

As built precedently, the local computational scheme may straightforwardly be implemented on a massively parallel computer. The principle is then extremely simple and efficient. Nevertheless, for practical reasons, implementation on conventional sequential computers is also discussed. In fact, most of the problems have already been solved by Vidale (1988, 1990), although on the basis of a radically different local computational scheme. In this section, for simplicity, the source will be considered located at a grid-point. Some generalizations are presented later, together with discussion of accuracy and initialization procedures.

### *Massively parallel implementation*

On a massively parallel computer, computation of first arrival times may be seen as a mere relaxation process. Initially, times are infinite everywhere, zero at the source point. Every grid-point is associated with a processor, so that during an elementary iteration, all grid-points simultaneously compute their arrival time according to the state of their nearest neighbours. This process is repeated until no further update is possible: all allocated times are then minimum. The number of iterations needed is of the same order as the maximum distance (in $h$ units) from source to model boundaries. The only precaution to be added to the local scheme is a test excluding the use of any difference between undefined (i.e., infinite) times during computations. This process was implemented in CM-FORTRAN on a Thinking Machines CM2 computer with 32k processors. Performances critically depend on optimization of data transfer between processors allocated to neighbour grid-points. On these grounds, standard CM-FORTRAN commands should be replaced by more efficient specialized procedures.

A similar process may easily be simulated on a sequential computer. Although this simulation is very inefficient, it demonstrates the feasibility of parallel implementation. In this simulation, computations are treated as asynchronous. During an elementary iteration, grid-points are examined in random order until one arrival time is changed (i.e., lowered). Computation ends when no arrival time was changed at the end of an iteration. Fig. 8 shows an example of random propagation of computations on a very contrasted velocity model. Figs 8(a), (b) and (c) show three successive steps of computations. Grid-points are plotted as soon as their allocated time becomes finite. Most grid-points are updated several times before first arrival is actually obtained. The shape of the computed zone roughly resembles wavefront shapes as finally obtained (Fig. 8d).

### *Sequential implementation*

Sequential implementation is based upon the procedures defined by Vidale (1988, 1990). Local computation is propagated according to a scheme which roughly follows the geometry of actual wavefront propagation in the model. This process tends to minimize the number of calls to the local computation function. In turn, this function is much shorter because the direction of propagation is known in advance, so that only three neighbours of the current point and three values of slowness are considered in 2-D, seven neighbours and four values of slowness in 3-D. Fig. 9 shows the simplified data structures used in the sequential implementation. A significant change with respect to Vidale's algorithm comes from the fact that computations are allowed to propagate towards the source region in certain circumstances.

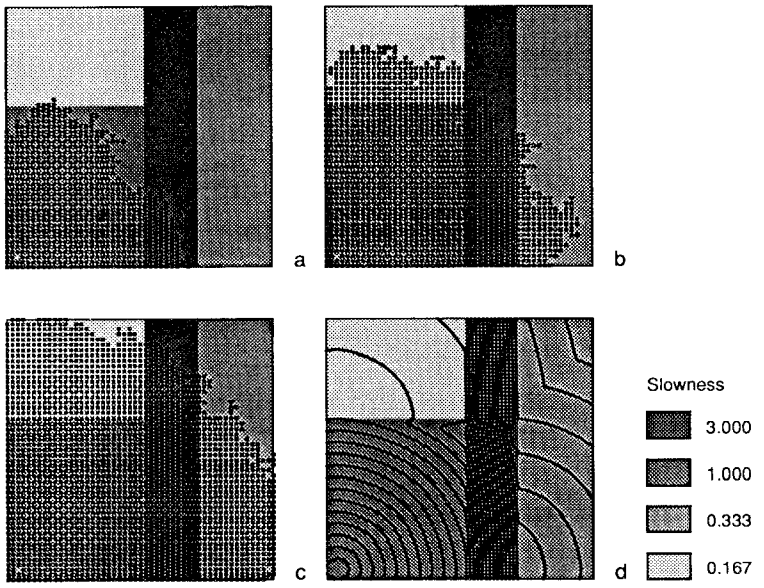


**Figure 8.** Three steps (a, b, c) of computations with the simulated parallel implementation. Points actually reached (finite arrival time) are indicated. Resulting time field (d) shows that order of computations roughly mimics propagating wavefronts.

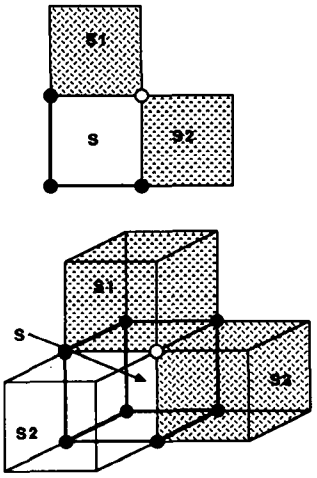


**Figure 9.** Simplified 2-D and 3-D data structures used in the sequential implementation. A hollow circle indicates the current point, black dots, its relevant neighbours. Values of slowness in the adjacent meshes are used for proper head wave generation along the current mesh boundaries.

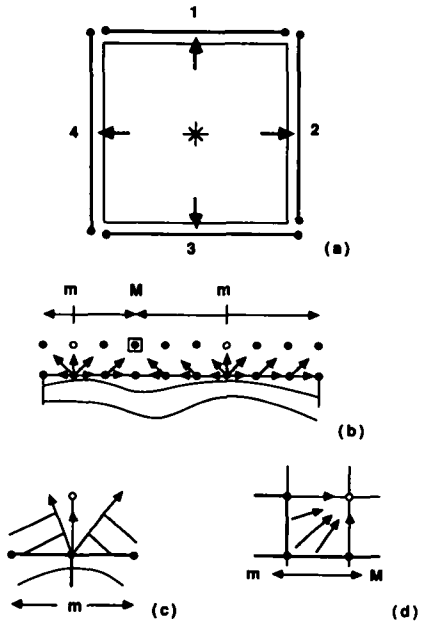


**Figure 10.** 2-D sequential implementation. (a) Sides of the current box are successively timed (note that corners are not treated separately, as in Vidale's algorithm). (b) Order of computations along a side. M and m indicate local time maxima and minima on the last computed parallel side. (c) Local wavefront pattern in front of a local time minimum. The current point is in a local shadow zone. (d) Time estimation at ordinary grid-points. At most five time estimates are evaluated (two transmissions submitted to illumination conditions, one diffraction, two interface waves).

Two main conditions constrain this implementation:

(i) causality (as introduced by Vidale 1988) must be respected: times must be computed in increasing order, because all stencils actually propagate wavefronts from past to future; and

(ii) no potential arrival time estimate may be ignored, so that an accurate first arrival is properly picked: every propagation mode (transmitted and diffracted body waves, head waves) must be examined.

*Implementation in 2-D.* In 2-D, computations normally proceed on square rings of increasing radius centred on the source. More complicated procedures designed to mimic wavefront geometry (Qin *et al.* 1990) are unnecessary and computationally costly. Each side of the current ring is successively examined (Fig. 10a). Arrival times are first computed in front of local time minima on the corresponding side of the previous ring. These points are those where the earliest wavefronts emerge from the previously computed zone. No particular stencil is necessary to deal with these points, as was the case in Vidale's method. On a local viewpoint, they systematically appear to be in a local shadow zone (diverging plane wavefronts, Fig. 10c), so that the point behind them (local time minimum) acts as a secondary source. As a consequence, arrival times at these points are simply computed by 1-D transmission.

Computations are then propagated from local time minima to local time maxima until all arrival times are

computed on the current side (Fig. 10b). All acceptable transmission and diffraction stencils are systematically computed. In particular, correct head wave generation involves the values of slowness in several meshes surrounding each grid-point. This local computation is explained in Fig. 10(d).

Nevertheless, two remaining problems must be faced when a head wave is generated along the computed side, which, in turn, may happen if this side contains an interface from a slow medium to a fast medium. First, this head wave may provide a first arrival at other points of the current side regardless of the timefield pattern on the previous side, whereas the order of computations defined earlier is actually based on this pattern. For this reason, when such a head wave is generated (i.e., the corresponding time estimate was found to be minimum), it must immediately be propagated from the current point to the extremity of the current side, so that this particular arrival is certainly timed at all relevant grid points. Finally, once the side has been entirely timed, it may be necessary to update times by propagating head waves into the computed zone, i.e., in the reverse direction. For this purpose, normal propagation has to be interrupted when the current side computation is completed. Reverse propagation then begins (Fig. 11): computations proceed on rows parallel to the last computed side, towards the computed zone, until it reaches a row where no arrival time is lowered. Normal (outwards) propagation is then resumed.

Reverse and normal propagations use exactly the same computational scheme because in zones where no head wave actually propagates, previously computed times will certainly not be lowered.

In some cases, reverse propagation itself is not sufficient to come to an end: back-propagated waves may in their turn generate interface waves so that twice reversed propagation must be allowed. In principle, this process should be seen as recursive. Reverse ($n$ times) propagation may provoke reverse ($n+1$ times) propagation. In practice, however, cases where $n>2$ may only be encountered in extreme situations in which the relevance of the slowness model should be questioned. Anyhow, the use of a programmation language allowing recursivity (e.g., C language) may deal with this problem in a very simple way.

This procedure avoids one of the main drawbacks in Vidale's algorithm where only outwards propagation was considered. As confirmation of this sequential implementation, tests showed that its results on a given model (for instance, the model used in Fig. 8) are bitwise identical to results obtained with the simulated parallel implementation.

*Implementation in 3-D.* In 3-D, computations proceed on cubic boxes. Principles defined in 2-D may easily be translated to 3-D. However, the order of computations on each square side of the current box appears less simple than in 2-D, because times computed on the preceding side may define a complicated topology, with hills and valleys. Causality implies that computations must always proceed 'uphill'. Although it is simple, Vidale's (1990) solution may still be made more straightforward. Causality is respected when points on the current side are computed in order of increasing times allocated to points behind them. The simplification arises from the fact that, as in 2-D (Fig. 10a), edges and corners of the current box need not be treated separately as in Vidale's algorithm.

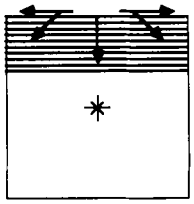

**Figure 11.** Reverse propagation. When head waves are generated along a side of the current box, arrival times may be updated (i.e., lowered) in the already computed zone. Reverse propagation proceeds on successive rows towards the source, until necessary.

As in 2-D, head wave generation along a side provokes, when its computation is completed, a shift from normal (outwards) propagation to reverse propagation from side to side towards the computed zone. As in 2-D, the process of reverse propagation is identical to normal propagation.

*Performance enhancements.* Performances of the sequential implementation depend on the structure of the slowness model. For instance, computational cost may be strongly penalized when many reverse propagation processes are generated. On these grounds, the case of a simple model with a strong vertical velocity gradient provides a perfect example. In such a model, first-order reverse propagation may be invoked so often that many points are updated up to hundreds of times before the actual first arrival is obtained. Performances may then be dramatically improved if computations proceed on rectangular boxes, so that the potential back-propagating sides (i.e., sides perpendicular to velocity gradient) are as short as possible. No general rule may easily be defined, as the algorithm may be applied to any velocity structure, but in some cases, very simple rules of thumb may provide spectacular performance enhancements.

## Initialization procedures

As local computations are based on a plane wavefront approximation, time estimates are quite imprecise at the vicinity of the source. Although in principle initialization might simply consist of setting zero arrival time at the source (or computing first arrivals at the closest grid-points by ray tracing), an appreciable gain in precision may be obtained at the expense of slightly more complicated initialization procedures.

When the source is surrounded with a constant velocity zone, direct arrivals should be computed exactly during initialization. A simple procedure may for instance determine the radius $r$ of the largest box in which slowness is constant and compute direct arrivals up to box $r-1$ (because box $r$ may contain back-propagating interfaces, it should be timed by finite differences). In practice, when $r>10h$, most of the imprecision generated by the plane wavefront approximation is avoided.

When the source is in heterogeneous zone, the finite difference scheme may be applied to a small zone around the source (e.g., up to radius $10h$) with a smaller mesh spacing (e.g., $h/2$). This process might even be used recursively. Imprecision is then concentrated into a smaller zone and its relative weight is strongly decreased.

The algorithm may also be modified in order to simulate other kinds of sources. For instance an incident plane

wavefront is easily simulated by computing conventional traveltimes on two borders of the model, according to the angle of incidence. This possibility was used for instance to obtain Fig. 7. As proposed by Moser (1989) other source configurations (e.g., the exploding reflector case) may also easily be implemented.

### Accuracy tests

The role of initialization is clearly demonstrated by accuracy tests.

Rather than comparing results obtained with the proposed algorithm with those provided by ray tracing techniques on complex heterogeneous models, we tested the accuracy of this scheme, in a first stage, against exact solutions on simple models. In a second stage, comparison with first arrival times picked on synthetic seismograms computed by full wave finite differences is presented.

The first model tested consists of a homogeneous layer lying on a homogeneous half-space with a velocity which is twice the velocity of the layer. Fig. 12(a) shows the wavefront pattern obtained by the FD computation, where transmitted and refracted waves are clearly identifiable. Fig. 12(b) is a plot of relative errors in per cent with respect to the exact solution obtained when initialization is restrained to setting zero arrival time at the source. Errors far from the source (inherited from initial wavefront curvature) are up to 0.5 per cent, whereas, close to the source, they reach several per cent. When direct arrivals are computed exactly in the initial homogeneous zone (Fig. 12c), maximal error on the same model drops down to only 0.15 per cent, which is acceptable for any practical purpose. Maximal absolute error represents less than 20 per cent, of the mean discrete time step ($h\langle s\rangle$) associated with model discretization.

The second model shows a smooth velocity variation. It provides, in addition to the question of initialization, an insight into the role of discretization of the slowness model. Fig. 13(a) shows wavefronts obtained in a model where velocity increases linearly with depth. The velocity ratio between top and bottom is 1:3. In Figs 13(b) and (c), relative errors on results are mapped for computations proceeding on $40 \times 40$ meshes and $200 \times 200$ meshes. In both cases, exact initialization was achieved up to radius $10h$ from the source, so that differences in precision only depend on the choice of mesh spacing. Fig. 13(b) shows that even a quite coarse grid (velocity contrasts up to 5 per cent from mesh to mesh) provides acceptable results.

Another method to assess the accuracy of first arrival times consists of a comparison with synthetic seismograms computed by a full wave finite difference algorithm. We chose the classical model of a corner embedded in a homogeneous medium studied by numerous authors (e.g., Kelly *et al.* 1976; Virieux 1986). Synthetic seismograms were computed with a second-order acoustic method with a complete heterogeneous formulation so that variations of velocities were correctly taken into account. Fig. 14 shows results in the case of a fast corner. First arrival times computed by the proposed algorithm were superimposed on seismograms chosen in order to record different kinds of waves (transmitted, refracted and diffracted). All seismograms show very good agreement with first traveltimes computed by FD.

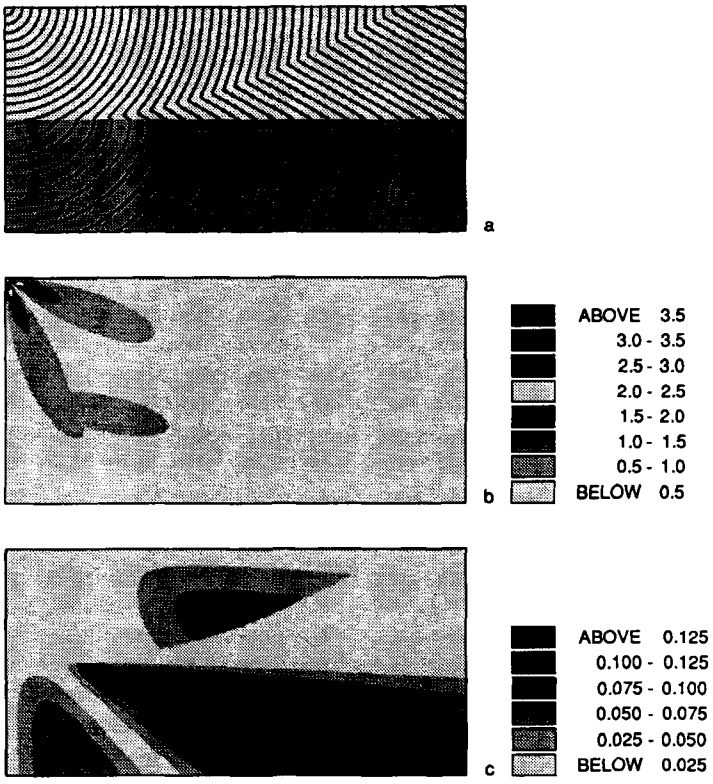


**Figure 12.** Comparison with exact solution. Velocity in the lower medium is twice as high as in the upper part. Model dimensions are $200 \times 100$ meshes. (a) Wavefronts. (b) Map of relative errors (in per cent) when initialization is restricted to the source point. (c) Errors obtained when exact direct arrivals are computed in the largest square homogeneous region surrounding the source.

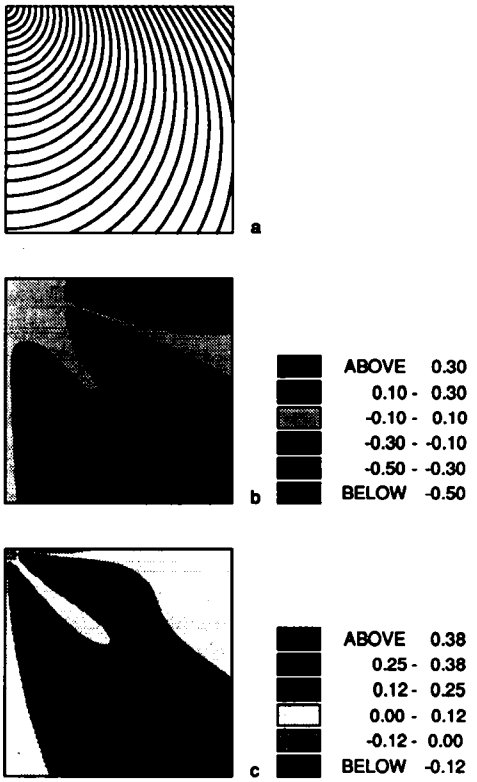


**Figure 13.** Comparison with exact solution. In this model, velocity increases linearly with depth by a factor 3. (a) Wavefronts. (b), (c) Maps of relative errors (in per cent) obtained when times are computed on $40 \times 40$ meshes (b) and $200 \times 200$ meshes (c). In both cases, exact initialization was achieved on $10 \times 10$ meshes.

Clearly, however, the plane wavefront approximation does not fail only at the vicinity of the source, but also at the vicinity of point diffractors (secondary sources) diagnosed during computations. However, an exhaustive discussion of accuracy would involve many independent factors (local imprecision due to wavefront curvature, azimuthal dispersion due to grid pattern, error propagation controlled by both local imprecision and integration path length). For this reason, general accuracy rules may not easily be defined, in order to securely propose good compromises between computational efficiency (memory requirements) and precision. On a purely heuristic basis, it appears that accuracy critically depends on the ratio between mesh spacing and mean distance between (primary or secondary) sources involved by computations. As a consequence, accuracy is definitely bad when mesh spacing is of the same order as dimensions of slowness anomalies in the model. In a sense, the rule is analogous with classical rules adopted in finite difference methods: mesh spacing should be small enough to avoid undersampling of the smallest spatial wavelengths of the slowness model. In principle, a systematic local rediscretization of the problem at the vicinity of diagnosed secondary sources might be included in the algorithm, but the computational cost of such a complication appears unnecessary.

## *A POSTERIORI* RAY TRACING

As proposed by Vidale (1988) and Moser (1989), rays may be 'backtracked' from any grid-point (receiver) to the source, with the use of the time field obtained by finite difference. Shooting rays from the source would be rather unpractical and, moreover, numerically unstable with respect to the initial ray parameter.

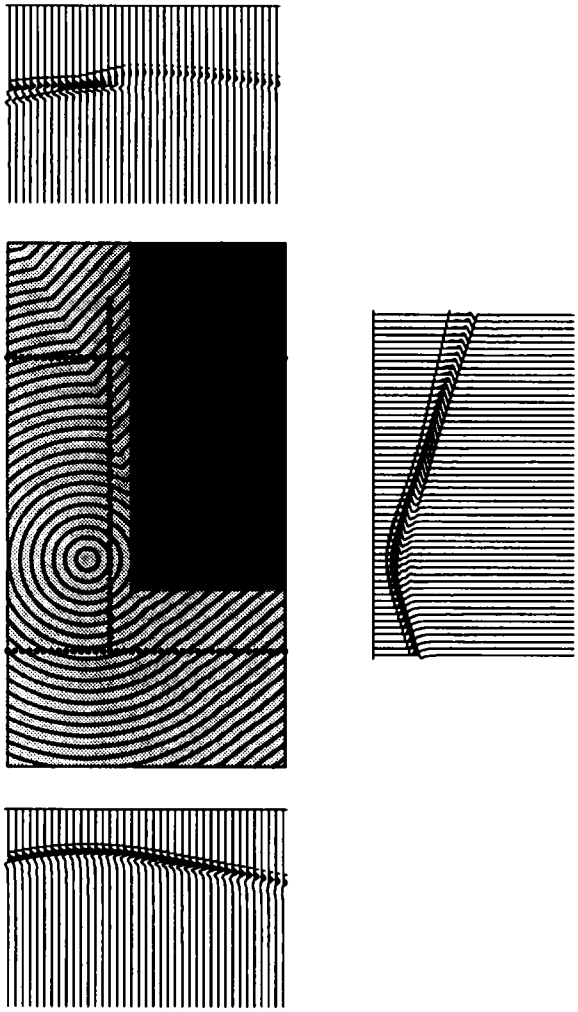

**Figure 14.** Comparison with full wave finite difference computations. Velocity in the upper right anomaly is 50 per cent higher than in the surrounding medium. Three synthetic seismic sections are plotted with superimposed ticks indicating computed first arrivals.

The ray connecting a given receiver to the source is the steepest path that may be found in the time field between these two points. Starting from the receiver, the ray is iteratively traced with increments opposite to the time gradient. For any point reached by the ray, the local time gradient is evaluated by a simple finite difference scheme at the closest grid-point or, if closer, at the centre of the mesh. The increment length is $h\sqrt{2}$, so that the ray 'ends' at a maximum distance $h$ from the source.

This possibility extends the uses of the first traveltime computation. For instance, in tomographic inversion procedures, ray trajectories are needed in order to convert time residuals into a localized update of the slowness model. *A posteriori* ray tracing may also be used as an accurate visual control of the relevance of first arrival time field. Even in the case of small errors on traveltime evaluation, irrelevant ray trajectories may be obtained, whereas wavefront patterns are much more robust with respect to small locations aberrations.

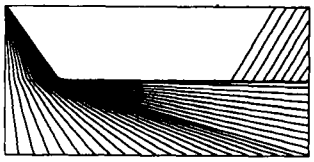

**Figure 15.** Rays traced in the model used in Fig. 12.

Figure 15 shows rays 'backtracked' from the borders of the grid, for the model used in Fig. 12. Rays refracted or transmitted according to Snell's law are easily identified. This figure shows moreover that these rays bear no information on a large part of the model in which the presence of any slow anomaly would not be detected with the exclusive use of first arrivals at this set of receivers.

## PERFORMANCE, EXTENSIONS, GEOPHYSICAL APPLICATIONS

In this section, we briefly discuss different uses and extensions of the proposed method. In particular, we stress questions which, when conveniently solved, could allow us to use a similar approach in order to compute signal amplitudes along with traveltimes. We also show that this method can be used to obtain reflected arrivals, so that it is not strictly limited to computation of first traveltimes.

### Performance in very contrasted models

Figure 16 illustrates how the proposed algorithm correctly takes into account very sharp first-order discontinuities in slowness, even in the case of complex topologies. The chosen model shows, in the upper region, a homogeneous disc embedded in a homogeneous medium in which the velocity is three times faster. An oblique interface separates this region from a still much faster medium. A source is located at the upper left corner of the model and wavefronts are plotted. Outside the disc, the effect of this slow anomaly is clearly shown by the appearance of diffracted wavefronts in the shadow zone. Inside the disc, the slowness contrast is such as to provoke a closure of wavefronts because rays travelling around the disc and refracted backwards are faster than any ray traced through the slow region. Head waves are also properly generated along the lower interface, although it is oblique with respect to coordinate axes. All interfaces are simply modelled as unsmoothed broken lines, in agreement with the discontinuous slowness model representation. The imprecision associated with this

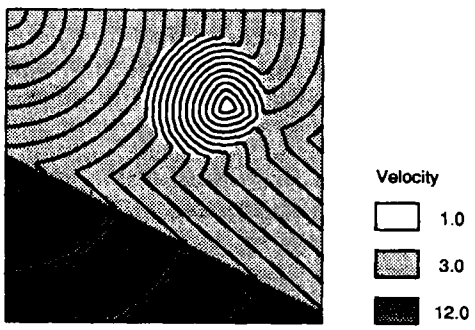


**Figure 16.** First arrival time field in a model with interfaces oblique with respect to coordinate axes. The grid size is 128 × 128.

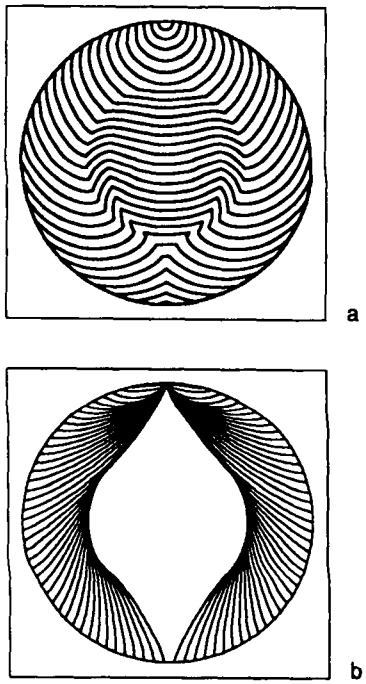


**Figure 17.** Computations in the PREM earth model, discretized over a Cartesian 560 × 560 grid (mesh spacing is 25 km). Slowness in the atmosphere was given a realistic value. (a) Wavefronts are plotted every 40 s. (b) Rays traced *a posteriori* from surface stations.

representation only depends on a reasonable choice of mesh spacing, i.e., of spatial resolution.

Although it is geologically irrelevant, this example shows that even very severe first-order contrasts raise no computational problem with the proposed method. This is also demonstrated in Fig. 17(a), where wavefronts are traced in the PREM earth model (Dziewonski & Anderson 1981). In this case, slowness in the atmosphere has been given a realistic value, so that very sharp contrasts occur both at the Earth's surface and at the core–mantle boundary.

### Evaluation of amplitudes: basic questions

Figure 17(b) shows rays traced in the earth model. Most of the rays are refracted at the core–mantle boundary. In terms of the classical ray approximation (i.e., at infinite frequency), such rays carry rigorously no energy. In the same way, comparison with synthetic seismograms in Fig. 14 shows that first arrivals may have such low amplitudes as to remain undetected in the presence of noise. For this reason, an evaluation of first arrival amplitudes would be helpful.

Vidale & Houston (1990) proposed a computation of amplitudes but they acknowledged their method does not correctly deal with sharp velocity contrasts. They geometrically estimate local wavefront curvature using differences between arrival times at a grid-point from four virtual sources (in the 2-D case) surrounding the actual source point. These differences are then converted to differences in take-off angle for neighbour grid-points. As in Vidale's algorithm, this approach does not take explicitly into account the presence of discontinuities of the time gradient,

so that mathematical errors may be encountered in these regions.

The underlying physical description of the problem may be summarized by a 'transport equation' implicitly dealing with *P*-waves (see, for instance, Jobert 1973):

$$\nabla(\rho A^2 \nabla t) = 0,$$

where $\rho(\mathbf{x})$ is the density and $A$ the amplitude of displacement. This equation may be seen as a conservation of energy along a ray tube of infinitely small radius. In principle, computation of amplitudes needs some more information (density field), but may then be treated as an ordinary differential problem. On this basis, a sensible improvement should be obtained if amplitudes are computed either along with arrival times, or afterwards, by finite differences.

This possibility is still under investigation, because the transport equation is an acceptable approximation only in regions where wavefront curvature is reasonable (Aki & Richards 1980; Ben Menahem & Beydoun 1985). In particular, this equation generates either zero or infinite amplitudes in regions where the time gradient is discontinuous. Those are the regions where Vidale & Houston's method also fails. Under these circumstances, the classical ray approximation must be abandoned and replaced by the generalized ray approximation involved by the geometrical theory of diffraction (Keller 1965; Felsen 1984). Once more, changing the mathematical approach of the problem does not reduce the underlying physical difficulties: computation of amplitudes with the ray tracing philosophy encounters exactly the same problems. With the finite difference approach, the most serious problem which remains to be solved concerns a realistic finite difference evaluation of amplitudes of head waves.

**Computation of reflected arrivals**

Van Trier & Symes (1991) claim that solutions of the eikonal equation picked by finite differences are generalized (viscosity) solutions which appear systematically to be first arrival time fields. Although they present this statement as an unproved mathematical conjecture, we should already consider it as a fundamental limitation to the finite difference approach. In fact, it is the main reason why the proposed method should certainly not be considered as a substitute to ray tracing.

Nevertheless, this conjecture does not state that finite difference computations may only provide first arrival times on a given model. We show here that, at the expense of some *a priori* information, several arrivals may be computed in some cases.

Figure 18(a) shows a schematic model in which the presence of a reflector is obvious. Nevertheless, this statement is definitely some supplementary information, because the finite difference scheme simply ignores reflection: reflected arrivals may never be first. The principle for computing reflected arrivals was explained by Williamson (1990), and also used by Aldridge & Oldenburg (1990) and Mtsuoka & Esaka (1990). In the sense of Huygens' principle, the picked reflector may be considered as a set a secondary sources. The reflected field is built as the result of the interference of wavelets emitted by this set of secondary sources. This is usually assumed as a conceptual basis for migration. The reciprocity principle applied to traveltimes implies that source and receiver play a symmetrical part (this would not be true if amplitudes were to be computed). For a given source–receiver pair, arrival times from the source and from the receiver to any point of the reflector are then first computed. These arrival times are added at each point of the reflector, indicating when individual wavelets are recorded at the receiver. Actual reflecting points (i.e., points where constructive interference patterns are obtained) are finally provided by Fermat's principle, which states that actual ray paths are those for which the traveltime from source to receiver is stationary.

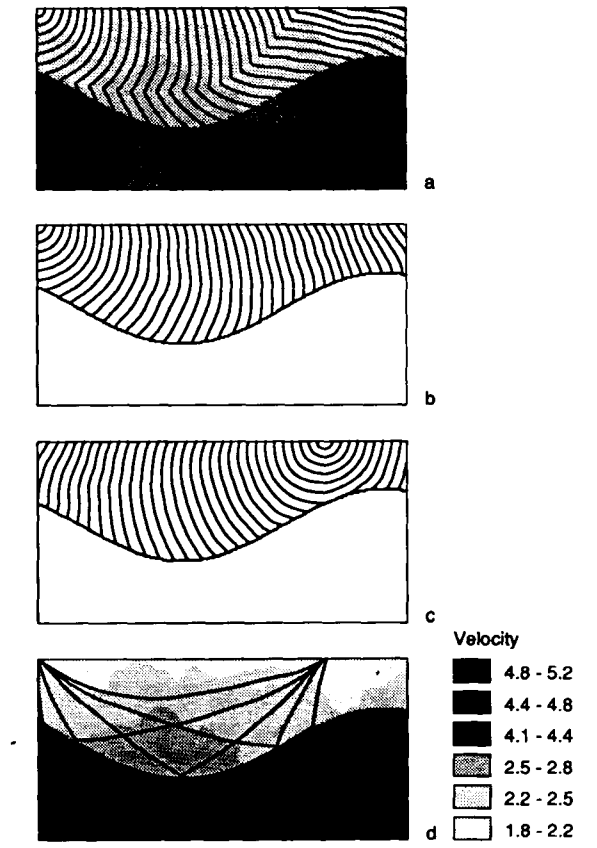


**Figure 18.** Computation of reflected arrivals. The model contains a (discretized) sine-shaped reflector. Both upper and lower regions show a vertical velocity gradient complicated by smooth perturbations (exponential noise). Mean velocity contrast at the reflector is 1:1.6. Grid size is 256 × 128. (a) First arrival time field obtained. (b), (c) Direct arrivals from source and receiver computed with an infinite slowness mask replacing the lower medium. Refracted wavefronts are avoided. (d) Rays traced after application of Fermat's principle (stationary arrival times).

Although this process is simple, some more precautions must be taken, because, as shown in Fig. 18(a), first arrivals from source to some part of the reflector may in fact be interface waves travelling along its surface with the velocity of the underlying medium. This may be the case whenever this underlying medium is faster than the upper one. Clearly, the associated upgoing head wave is irrelevant to the reflection problem, which should only involve the slowness in the upper part of the model. In other terms, we must compute *direct* arrivals from source (and receiver) to reflector, although some of them may not be first arrivals.

This is very simply achieved by computing first arrivals in a different model in which the underlying medium is replaced by a zero velocity medium. This infinite slowness mask prevents the generation of interface waves along the reflector and ensures that computed times only depend on the overlying slowness model. Figs 18(b) and (c) show wavefronts obtained with this masked computation. The three reflected Fermat rays obtained are represented in Fig. 18(d) for visualization of the result, but no ray tracing is actually involved in this method.

In this particular case, up to five different arrivals may thus be computed at the receiver: refracted (first), direct and three reflected events. It should be noted that nothing implies that these are the five first events recorded at the receiver, because, for example, some other non-negligible reflector may have been overlooked in the upper region of the model.

For a set of $n$ sources and receivers this method involves $n$ first traveltimes computations (which may easily be achieved in parallel), so that the computational cost is reasonable. It allows computation of reflected arrival times in very complex slowness models containing virtually any reflector morphology. Moreover, with the use of infinite slowness masks, the same approach might be used in order to obtain a series of refracted and reflected arrivals in models consisting of a stack of layers. No constraint would then be imposed on the complexity of the slowness model adopted between two successive reflectors.

### Geophysical applications

In his first paper, Vidale (1988) already pointed out most of the various potential applications of a very fast first traveltime computation method. The combination of FD traveltime computation and *a posteriori* ray tracing makes the proposed method of potential tool for many geophysical applications, some of them poorly managed with standard ray tracing techniques.

For instance, topographic corrections for the effect of sea-floor or Earth's surface topography remain a largely open question with the ray tracing approach, although numerous authors proposed various correction procedures (e.g., Purdy 1982, for marine data). The robustness of the FD algorithm with respect to sharp arbitrarily shaped velocity contrasts makes it possible to introduce actual topographic data in the model used and straightforwardly adopt the velocity of sound in sea-water for marine studies, or in air for land studies. It should be repeated here, however, that precise computations demand that mesh spacing be smaller than the smallest spatial wavelengths of the topographic data used.

In the same way, this method might provide new insights on first traveltime inversion, allowing inversion schemes to use both global information (wavefront patterns) and localized data (ray paths). Moreover, thanks to computational speed, inversion with non-linear techniques such as simulated annealing might be investigated on a realistic basis.

Up to now, Vidale's approach was only applied to earthquake location (Vidale 1988; Ammon & Vidale 1989; Nelson & Vidale 1990). A significant improvement with respect to conventional localization methods comes from the fact that, with the help of the recriprocity principle, arrival times at a given set of receivers may be very quickly translated, for a given velocity model, into a global cost function map which may be interpreted as a test of model accuracy: imprecision on source localization is due both to noise (imprecision on picked times) and model misfits. The associated ray fields may then be used to update the model iteratively.

## CONCLUSION

Some other applications already proposed by Vidale (1988), such as migration in arbitrary velocity models, should also be investigated. Clearly, although finite difference computation of first arrival times is not intended to rule out the ray tracing approach, a large spectrum of applications is already well identified, and some more might appear in the near future. Two main improvements are provided by the proposed method, in comparison with previous algorithms, i.e., robustness with respect to very sharp velocity contrasts is warranted, and an extremely fast massively parallel implementation is defined. Moreover, introduction of carefully designed infinite slowness masks offers ways to compute several arrival times with the same scheme.

## ACKNOWLEDGMENTS

We wish to thank J. Virieux for his helpful comments. The paper benefited from stimulating discussions with V. Farra, D. Gibert and G. Wittlinger. A constructive review was made by J. Vidale. The Geophysical Research group at Ecole des Mines is supported by CGG, CFP, ELF and IFP. One of the authors (IL) was supported for a PhD thesis by Institut Français de Recherche pour l'Exploitation de la Mer (IFREMER), France. Parallel computations tests were achieved at Institut de Physique du Globe de Paris on a CM2 Connexion Machine.

## REFERENCES

Aki, K. & Richards, P. G., 1980. *Quantitative Seismology: Theory and Methods*, W. H. Freeman and Co., San Francisco.

Aldridge, D. F. & Oldenburg, D. W., 1990. Refractor imaging using an automated wavefront reconstruction method, *Soc. Exploration Geophys. 1990 Meeting*, Expanded abstracts, pp. 1575–1577, SEG, Tulsa, OK.

Ammon, C. J. & Vidale, J., 1990. Seismic traveltime tomography using combinatorial optimization techniques (abstract), *Seism. Res. Lett.*, **61**, 39.

Ben Menahem, A. & Beydoun, W. B., 1985. Range of validity of seismic ray and beam methods in general inhomogeneous media. I. General Theory, *Geophys. J. R. astr. Soc.*, **82**, 207–234.

Dziewonski, A. M. & Anderson, D. L., 1981. Preliminary Reference Earth Model, *Phys. Earth planet. Inter.*, **25**, 297–356.

Felsen, L. B., 1984. Geometrical theory of diffraction, evanescent waves, complex rays and Gaussian beams, *Geophys. J. R. astr. Soc.*, **79**, 77–88.

Jobert, G., 1973. Propagation des ondes en milieux homogèges, in *Traité de Géophysique Interne*, vol. 1, pp. 107–109, eds Coulomb, J. & Jobert, G., Masson, Paris.

Keller, B., 1965. Geometrical theory of diffraction, *J. Opt. Soc. Am.*, **52**, 116–130.

Kelly, K. R., Ward, R. W., Treitel, S. & Alford, R. M., 1976. Synethic seismograms: a finite difference approach, *Geophysics,* **41,** 2–27.

Moser, T. J., 1989. Efficient seismic ray tracing using graph theory, *Soc. Exploration Geophys. 1990 Meeting,* Expanded abstracts, pp. 1106–1108, SEG, Tulsa, OK.

Mtsuoka, T. & Ezaka, T., 1990. Ray tracing on reciprocity, *Soc. Exploration Geophys. 1990 Meeting,* Expanded abstracts, pp. 1028–1031, SEG, Tulsa, OK.

Nelson, G. D. & Vidale, J., 1990. Earthquake location by 3D finite differences traveltimes, *Bull seism. Soc. Am.,* **80,** 395–410.

Purdy, G. M., 1982. The correction for the traveltime effects of seafloor topography in the interpretation of marine seismic data, *J. geophys. Res.,* **87,** B10, 8389–8396.

Qin, F., Olsen, K. B., Luo, Y. & Schuster, G. T., 1990. Solution of the eikonal equation by a finite difference method, *Soc. Exploration Geophys. 1990 Meeting,* Expanded Abstracts, pp. 1004–1007, SEG, Tulsa, OK.

Riznichenko, Yu. V., 1946. Geometrical seismics of layered media. Trudy Inst. Theor. Geophysics, Vol. II, *Izd. An. SSSR,* Moscow (in Russian).

Saito, H., 1989. Traveltimes and raypaths of first arrival seismic waves: computation method based on Huygens' principle, *Soc. Exploration Geophys. 1989 Meeting,* Expanded abstracts, pp. 244–247, SEG, Tulsa, OK.

Saito, H., 1990. 3-D ray tracing method based on Huygens' principle, *Soc. Exploration Geophys. 1990 Meeting,* Expanded abstracts, pp. 1024–1027, SEG, Tulsa, OK.

Thornburgh, H. R., 1930. Wave-front diagrams in seismic interpretation, *AAPG Bull.,* **14,** 185–200.

Van Trier, J. & Symes, W. W., 1990a. Upwind finite-difference calculation of seismic traveltimes, *Soc. Exploration Geophys. 1990 Meeting,* Expanded abstracts, pp. 1000–1002, SEG, Tulsa, OK.

Van Trier, J. & Symes, W. W., 1991. Upwind finite-difference calculation of traveltimes, *Geophysics,* in press.

Vidale, J., 1988. Finite-difference calculation of travel times, *Bull. seism. Soc. Am.,* **78,** 2062–2076.

Vidale, J., 1990. Finite-difference calculation of travel times in 3D, *Geophysics,* **55,** 521–526.

Vidale, J. & Houston, H., 1990. Rapid calculations of seismic amplitudes, *Geophysics,* **55,** 1504–1507.

Virieux, J., 1986. P–SV wave propagation in heterogeneous media: velocity–stress finite difference method, *Geophysics,* **51,** 889–901.

Williamson, P. R., 1990. Tomographic inversion in reflection seismology, *Geophys. J. Int.,* **100,** 255–274.