STABLE ITERATIVE RECONSTRUCTION ALGORITHM FOR NONLINEAR TRAVELTIME TOMOGRAPHY

James G. Berryman
Lawrence Livermore National Laboratory
University of California
P. O. Box 808 L-156
Livermore, CA 94550

Short title: Nonlinear Traveltime Tomography

Classification numbers: 91.30 43.63 87.60B 42.40

ABSTRACT

Reconstruction of acoustic, seismic, or electromagnetic wave speed distribution from first arrival traveltime data is the goal of traveltime tomography. The reconstruction problem is nonlinear, because the ray paths that should be used for tomographic backprojection techniques can depend strongly on the unknown wave speeds. In our analysis, Fermat's principle is used to show that trial wave speed models which produce *any* ray paths with traveltime smaller than the measured traveltime are *not* feasible models. Furthermore, for a given set of trial ray paths, nonfeasible models can be classified by their total number of "feasibility violations", *i.e.*, the number of ray paths with traveltime less than that measured. Fermat's principle is subsequently used to convexify the fully nonlinear traveltime tomography problem. In principle, traveltime tomography could be accomplished by solving a multidimensional nonlinear constrained optimization problem based on counting the number of ray paths that exactly satisfy the measured traveltime data. In practice, this approach would be too computationally intensive without the use of massive parallel computing architecture. Nevertheless, the insight gained from from this new point of view leads to a stable iterative reconstruction algorithm. The new algorithm is a modified version of damped least-squares (also known as "ridge regression"). The correction step at each iteration is in the direction of the damped least-squares solution, but the size of the step is determined by the location of the point having the minimum number of feasibility violations in the direction of the step. The computational burden of computing the number of feasibility violations is virtually negligible. Examples of the results produced by this algorithm are given.

1. Introduction

The classical methods of tomography provide a means for reconstructing a two-dimensional function from a set of line integrals [1,2]. For medical x-ray tomography [3], such line integrals are provided by measurements of the wave amplitude attenuation for straight rays passing through the body. When back-projected along the known ray paths [4,5], the attenuation data provide a picture of an inhomogeneous density distribution; the picture can then be interpreted for purposes of diagnosis.

Tomography has many uses outside of the field of medicine – including electron microscopy [1,2,4], acoustical and optical tomography [2,5], and radio astronomy [6]. In geophysical applications to the whole earth [7,8] or to local reconstruction problems such as borehole-to-borehole scanning with either electromagnetic or seismic probes [9,10], the assumption of straight ray paths is often a poor approximation [11,12]. Nevertheless, sophisticated full wave field processing schemes such as geophysical diffraction tomography [13] (using both wave amplitude and phase in the reconstruction) are known to work very well when a starting model of sufficient accuracy is available. Thus, there is reason to believe that the reconstruction problem with ray bending might be solved if some sensible procedure for finding a good starting model (other than trial and error) could be found. Recent progress towards finding an iterative algorithm for computing just such a model from traveltime data is described in this paper.

Traditional traveltime tomography reconstructs a slowness (reciprocal wave speed) model from measured traveltimes for first arrivals. The locations of sources and receivers are assumed known, but the actual ray paths are not known and must be determined along with the model slowness (see Figure 1). Fermat's principle [14] – that the path taken is the one of least traveltime – has been used extensively in forward modeling, *i.e.,* given the slowness model Fermat's principle determines the ray paths. However, Fermat's principle may also be applied in an entirely different way during the reconstruction of the slowness model using traveltime data [15], as we shall show.

The main purpose of this paper is to develop a stable method of reconstructing wave speed structure from first arrival traveltime data. Our emphasis is on problems containing high contrast anomalies assuming that *the errors made if one neglects ray bending effects are far more significant* than are those made by neglecting measurement errors in the traveltime data. Sections 2 and 3 develop notation and a geometrical description of the analysis that will make it easier to understand the main result of the paper presented in section 4. In section 2, we show how Fermat's principle may be used to convexify the fully nonlinear traveltime tomography problem. It is also shown that, for a fixed ray-path matrix, the total number of "feasibility violations" may be used to classify all wave speed models in the nonfeasible region. In section 3, relaxed constraints are introduced and their relationship to weighted least-squares solutions is found. In section 4, the number of feasibility violations is used to choose the magnitude of the correction step in each iteration of our modified ridge regression method. Numerical results are presented in section 5. Our conclusions are summarized in section 6.

2. Fermat feasibility boundary

To set notation, let t be the measured traveltime m-vector such that $t^T = (t_1, \ldots, t_m)$, where t_i is the traveltime along the i-th ray path (a superscript T implies the transpose) and m is the total number of views (transmitter-receiver pairs). We form our model in two-dimensions (see Figure 2) by dividing the rectangular region enclosed by our sources and receivers into rectangular cells of constant slowness. In three dimensions, the cells are blocks of constant slowness. Then, s is the model slowness n-vector $s^T = (s_1, \ldots, s_n)$, with s_j being the slowness of the j-th cell and n the total number of cells. For forward modeling, s and t are related by the equation

$$Ms = t, \tag{1}$$

where M is an $m \times n$ matrix whose matrix elements $l_{i,j}$ are determined by the length of the i-th ray path as it passes through the j-th cell. Eq. (1) simply states that the total traveltime along a ray path is the sum of the traveltimes through each of the cells traversed by the ray. Fermat's principle is often used in forward modeling to determine M and therefore t when s is given.

2.1 Feasibility constraints

The inverse problem associated with (1) starts with traveltime data t and attempts to find the corresponding slowness model s and ray-path matrix M. We will now depart from traditional methods by applying Fermat's principle in a new way [15]. The forward problem (1) is replaced by the m feasibility constraints

$$(Ms)_i \geq t_i. \tag{2}$$

This fact follows from Fermat's principle: the first arrival necessarily followed the path of minimum traveltime for the model s. Thus, (2) must be satisfied by *any ray-path matrix* M if s is the true model and therefore any model that violates (2) along any of the ray paths is not a feasible model. An exact solution to the inversion problem is found if and only if all of the inequalities in (2) become identities for some choice of model slowness vector s. (The question of uniqueness of the solution is discussed in subsection 2.4.) For each of the m inequality constraints (2), the limiting equality is the equation for a hyperplane in the n-dimensional slowness model space. The feasible region is bounded by these hyperplanes and by the planes determined by positivity of slowness in all cells j,

$$s_j > 0. \tag{3}$$

The two sets of inequalities (2) and (3) guarantee that the feasible region of the model space is convex. Thus, for fixed ray-path matrix M, the set of all feasible models s includes all models either inside the feasible region or on the feasibility boundary determined by M and t.

So far the argument has been pertinent only to linear traveltime tomography (*i.e.*, fixed ray-path matrix M). However, it is a small step (see Figure 3) to see that the constraints (2) imply the existence of a definite convex set in the model space containing *all* the feasible models for arbitrary choices of the ray-path matrix: Since any point s that is nonfeasible for *any particular choice of* M must lie outside of the global feasibility set, it follows that the intersection of the feasible sets for all choices of M determines the global (nonlinear) feasibility set. This global set must be convex since it is the intersection of convex sets. Furthermore, an exact solution of the inverse problem (*i.e.*, assuming the data are consistent so such a solution exists) must lie on the boundary of this global convex set. Finally, we note that the location of the global feasibility boundary depends only on the set of measured traveltimes t, and on the locations of the transmitters and receivers.

2.2 Feasibility violation number

Another new concept that is useful in computations is that of "feasibility violation number" $N_M(s)$. For any combination of ray-path matrix M, slowness vector s, and measured traveltimes t, the number of rays violating the constraints (2) is given by

$$N_M(s) \equiv \sum_{i=1}^{m} \theta[t_i - (Ms)_i], \tag{4}$$

where the step function $\theta(x)$ is defined by

$$\theta(x) = \begin{cases} 0, & \text{for } x \le 0; \\ 1, & \text{for } x > 0. \end{cases} \tag{5}$$

The number $N_M(s)$ is equal to zero in the feasible region. Furthermore, it is clearly a monotonically increasing function of distance from the local feasibility boundary associated with M – once one of the hyperplanes of (2) is crossed we will never cross it again if we keep moving in the same direction in the model space. Thus, $N_M(s)$ is cheap to compute and gives us a rough idea of how close we are to the feasibility boundary.

The feasibility violation number may be used to classify all models in the nonfeasible region of the model space. For example, contours of constant N_M may be drawn starting with $N_M = 0$ on the feasibility boundary.

2.3 Scale invariance of M

Another important idea that will be used repeatedly in the sections that follow is that of scale invariance. In forward modeling, we use Fermat's principle to find the optimum ray-path matrix M associated with a given slowness model s. However, the M found this way is actually optimum for all slowness models in the same direction as s, *i.e.*, all models of the form γs where γ is an arbitrary positive scalar:

$$\gamma t_i = \min_{\{paths\}} \int_{path} \gamma s dl_i.$$

Thus, the optimum M is associated with a direction, not just with a single point, in the model space. This scale invariance property of the ray-path matrices is very useful in the reconstruction algorithms, because it means that any model can be scaled to achieve either feasibility or some other desirable property (such as an improved fit to the traveltime data) without affecting the validity of the current ray-path matrix M.

2.4 Lack of uniqueness

There are three principal sources of nonuniqueness in traveltime tomography: (*i*) choice of model (slowness) parameterization, (*ii*) measurement errors, and (*iii*) ghosts.

(*i*) The first type of nonuniqueness arises from our initial (somewhat arbitrary) choice of the parameterization of the slowness model. In real problems, the medium to be analyzed may be assumed to have an essentially continuous distribution of slowness. So any particular discrete parameterization of such a medium is subject to the criticism that it is not providing a realistic model of the medium. However, as long as the wavelengths of the acoustic, seismic, or electromagnetic probe are comparable to or greater than the cell size, some sort of discrete approximation is entirely justified. Even if the model does not have the resolution we would like it to have, it may still be entirely consistent with the data that is collected. After one accepts a discrete parameterization, there still remains the nonuniqueness associated with our particular choice of cell size. The question of optimum cell size will not be treated here, but the question can be addressed in a straightforward (if *ad hoc*) manner by trying different parameterizations for the same data. If the results are qualitatively insensitive to the parameterization, we may have some confidence in their validity. Having said this, we will not treat this issue further in the present paper.

(*ii*) The second type of nonuniqueness arises from any measurement errors in the traveltime data (whether real or synthetic). These errors may be due to errors in transmitter or receiver location, or they may be due to errors in correctly picking the first arrivals and/or measuring the absolute time from transmission to reception of the signal. Even synthetic data may be expected to have some errors associated with the ray tracing or finite difference algorithm used to produce the data. The consequences of such errors on the nonuniqueness of the reconstruction are similar to those arising from the choice of model parameterization. Our method of treating them is corresponding similar. We can check the sensitivity of the final results to such errors by varying these data values within the expected range and then determine how sensitive the reconstructed values are to these changes. Lack of sensitivity to small errors is a sign of a robust reconstruction method. Later in the paper we will treat the effects of measurement errors further.

(*iii*) Once having chosen the model parameterization and having neglected the measurement errors, the third source of nonuniqueness is still difficult to treat. It arises because uniqueness of the solution s is possible only when the matrix M is of full rank — which is highly unlikely in nonlinear traveltime tomography. The subject of such ghosts [16,17] in tomography is well-known, a ghost being any slowness vector in the null-space. Thus, any n-vector σ from the null space of M (*i.e.*, $M\sigma = 0$) may be added to a solution s satisfying $Ms = t$ without affecting the agreement with the traveltime data through $M(s+\sigma) = t$. As a simple example, suppose that none of the rays making up the ray-path matrix pass through cell k. Then, the value of the slowness s_k is arbitrary and any value may be added to the slowness of this cell. We should therefore think of the solution s that we seek as a particular solution while any vector σ from the null space is a homogeneous solution. The full solution of our problem must include a careful analysis of the span of the null space of the ray-path matrix M associated with the particular solution s. Clearly, the primary goal of the analysis must be to find such a particular solution and its associated ray-path matrix M. The analysis of the corresponding null space — which may be carried out using standard numerical techniques (*i.e.*, singular value decomposition) — is a secondary goal and will not be a major concern for us in the present paper.

3. Significance of relaxed constraints

In principle, traveltime tomography could be done by solving a multidimensional nonlinear constrained optimization problem based on counting the number of ray paths that exactly satisfy the measured traveltime data, *i.e.*, moving along the feasibility boundary using something akin to a simplex method [18]. In practice, this approach would be too computationally intensive for most applications; the global feasibility boundary is determined only implicitly by (2) and, because of the high dimensionality of the model space, would require a very large number of forward calculations for different Ms to produce a useful map of this boundary. So we will seek other ways of using the insight gained from this new point of view to produce a stable iterative reconstruction algorithm. In particular, we will permit feasibility violations during the reconstruction since the nonfeasible region has the most easily quantified structure. We can always use the scale invariance property of M to place the model slowness on the feasibility boundary at the end of the calculation if desired.

3.1 Relaxed constraint

Let $u^T = (1, \ldots, 1)$ be an m-vector of ones. Then, the total traveltime along all the ray paths is given by $u^T t$. One method of relaxing the constraints (2) is to sum over these inequalities to produce a single constraint

$$u^T M s \geq u^T t. \tag{6}$$

The limiting equality in (6) determines a single hyperplane constraint. This hyperplane has the physical significance of being the unique hyperplane associated with the ray-path matrix M for which all slownesses s have the same total traveltime as that measured. All models feasible for M satisfy (6), but some nonfeasible models also satisfy (6). That is why (6) is called a relaxed constraint.

3.2 Tailored eigenvalue problem

There are at least two important eigenvalue problems associated with the ray-path matrix M. Following Lanczos [19], we will introduce the $(m+n)\times(m+n)$ real, symmetric matrix H determined by M

$$H = \begin{pmatrix} 0 & M \\ M^T & 0 \end{pmatrix}. \tag{7}$$

Then, the first eigenvalue problem has the form

$$\begin{pmatrix} 0 & M \\ M^T & 0 \end{pmatrix} \begin{pmatrix} u_\lambda \\ v_\lambda \end{pmatrix} = \lambda \begin{pmatrix} L & 0 \\ 0 & C \end{pmatrix} \begin{pmatrix} u_\lambda \\ v_\lambda \end{pmatrix}. \tag{8}$$

In (8), the vectors u_λ and v_λ are of length m and n respectively. The matrix on the right is defined in terms of the diagonal matrices L and C whose diagonal elements are the row sums L_i and column sums C_j

$$L_i = \sum_{j=1}^{n} l_{i,j}, \qquad C_j = \sum_{i=1}^{m} l_{i,j}. \tag{9}$$

The quantity L_i is seen to be the total length of path i. The quantity C_j is the total length of all the ray-path segments that pass through cell j, so we will call this the "coverage" of cell j. Any cell with $C_j = 0$ is uncovered and therefore lies outside the span of our data for the current choice of ray paths. We retain only the covered cells in the reduced slowness vector $\tilde{s}$ of length $\tilde{n} \le n$. The matrix M may similarly be reduced to $\widetilde{M}$ by deleting the corresponding columns of zeros. Finally, the diagonal matrix C is modified to include only the nonzero sums in (9). For simplicity of notation, we assume that $\tilde{n} = n$ in the following discussion.

Now recall $u^T = (1, \ldots, 1)$, an m-vector of ones, and define $v^T = (1, \ldots, 1)$, an n-vector of ones. Then,

$$Mv = Lu \tag{10}$$

and

$$M^T u = Cv. \tag{11}$$

Problem (8) has been treated in detail by Berryman [20], and is most important for reconstruction problems having low to moderate contrasts ($< 20\%$). We see in particular that Eqs. (10) and (11) are of the form (8) so that the $(m+n)$-vector of ones $(u^T, v^T) = (1, \ldots, 1)$ is an eigenvector of H with eigenvalue unity. This fact is significant, and provides the motivation for studying the eigenvalue problem (8), as has been shown elsewhere [20].

An analogous eigenvalue problem pertinent for high contrast reconstructions takes the form

$$\begin{pmatrix} 0 & M \\ M^T & 0 \end{pmatrix} \begin{pmatrix} w_\lambda \\ x_\lambda \end{pmatrix} = \lambda \begin{pmatrix} \hat{T} & 0 \\ 0 & \hat{D} \end{pmatrix} \begin{pmatrix} w_\lambda \\ x_\lambda \end{pmatrix}. \tag{12}$$

where, for $\lambda = 1$, $w_1 = u$, and $x_1 = \hat{s}$ is the current best estimate of the reconstructed slowness. Then, by analogy with (10) and (11),

$$M\hat{s} = \hat{T}u, \tag{13}$$

$$M^T u = Cv = \hat{D}\hat{s} \tag{14}$$

where $\hat{T}$ and $\hat{D}$ are diagonal matrices with elements given by

$$\hat{T}_i = \sum_{j=1}^{n} l_{i,j}\hat{s}_j \qquad \text{for } 1 \le i \le m \tag{15}$$

and

$$\hat{D}_j = \sum_{i=1}^{m} l_{i,j}/\hat{s}_j = C_j/\hat{s}_j \qquad \text{for } 1 \le j \le n. \tag{16}$$

The $\hat{T}_i$s are the traveltimes of the current ray paths through the current model $\hat{s}$. The $\hat{D}_j$ is the coverage of the j-th cell divided by the slowness of that cell; the dimensions of $\hat{D}_j$ are (length)^2/time — or, the same as that of a diffusion coefficient. We will call $\hat{D}_j$ either the "modified cell coverage" or the "cell diffusion factor."

Now we will transform (12) to a canonical form using

$$\begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix} = \begin{pmatrix} \hat{T}^{-\frac{1}{2}} & 0 \\ 0 & \hat{D}^{-\frac{1}{2}} \end{pmatrix} \begin{pmatrix} 0 & M \\ M^T & 0 \end{pmatrix} \begin{pmatrix} \hat{T}^{-\frac{1}{2}} & 0 \\ 0 & \hat{D}^{-\frac{1}{2}} \end{pmatrix} \tag{17}$$

so that

$$\begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix} = \begin{pmatrix} 0 & \hat{T}^{-\frac{1}{2}} M \hat{D}^{-\frac{1}{2}} \\ \hat{D}^{-\frac{1}{2}} M^T \hat{T}^{-\frac{1}{2}} & 0 \end{pmatrix} \tag{18}$$

and

$$\begin{pmatrix} y_\lambda \\ z_\lambda \end{pmatrix} = \begin{pmatrix} \hat{T}^{\frac{1}{2}} w_\lambda \\ \hat{D}^{\frac{1}{2}} x_\lambda \end{pmatrix}. \tag{19}$$

The motivation for using this particular choice of transform will be clarified in section 3.4. The eigenvalue problem (12) is then transformed into

$$\begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix} \begin{pmatrix} y_\lambda \\ z_\lambda \end{pmatrix} = \lambda \begin{pmatrix} y_\lambda \\ z_\lambda \end{pmatrix}. \tag{20}$$

The significance of (20) will become clear in the following discussion. We wish to emphasize now that, with the normalization that has been performed to produce A, *the current slowness model* $\hat{s}$ *gives rise to the unique eigenvector of (20) with the highest eigenvalue and that eigenvalue is unity — i.e.,* $A^TA\hat{z} = \hat{z}$.

3.3 Weighted least-squares

Now we will consider two weighted least-squares fitting problems. In both examples, the weights $\hat{T}$ and $\hat{D}$ have been incorporated directly into the normalization (preconditioning) factors for the matrices A, and the eigenvector components y_λ and z_λ. The first problem is to find the slowness $s = \gamma\hat{s}$ in the direction of $\hat{s}$ giving the best least-squares fit to the measured traveltime data. If the normalized traveltime measurement vector is given by

$$\hat{y} = \hat{T}^{-\frac{1}{2}}t, \tag{21}$$

then the problem is to find γ such that

$$\psi(\gamma) = (\hat{y} - A\gamma\hat{z})^T(\hat{y} - A\gamma\hat{z}) \tag{22}$$

achieves its minimum. This value is found to be

$$\gamma = \frac{\hat{z}^TA^T\hat{y}}{\hat{z}^TA^TA\hat{z}} = \frac{\hat{z}^TA^T\hat{y}}{\hat{z}^T\hat{z}}. \tag{23}$$

since $A^TA\hat{z} = \hat{z}$ by the constructions of subsection 3.2. If the optimum scale factor has already been found and the value $\hat{z}$ scaled appropriately, then $\gamma = 1$ in (23). We will assume this is the case for the remainder of this discussion.

Now consider a second weighted least-squares problem. Consider the objective function

$$\phi_\mu(z) = (\hat{y} - Az)^T(\hat{y} - Az) + \mu(z - \hat{z})^T(z - \hat{z}) \tag{24}$$

where μ is a damping parameter [21,22]. The minimum of (24) occurs at $z = z_\mu$ where z_μ satisfies

$$(A^TA + \mu I)(z_\mu - \hat{z}) = A^T\hat{y} - \hat{z}. \tag{25}$$

To arrive at (25), we again used the fact that $A^TA\hat{z} = \hat{z}$. The solution z_μ is called the damped least-squares solution or the ridge regression solution.

Now notice that the right hand side of (25) is orthogonal to $\hat{z}$, *i.e.*,

$$\hat{z}^T(A^T\hat{y} - \hat{z}) = 0. \tag{26}$$

Eq. (26) follows from (23) when $\gamma = 1$. Applying $\hat{z}^T$ to (25) then gives

$$(1+\mu)\hat{z}^T(z_\mu - \hat{z}) = 0, \tag{27}$$

so that z_μ lies in a hyperplane orthogonal to $\hat{z}$.

3.4 Hyperplane of constant total traveltime

To establish the significance of the relaxed constraints, now note that

$$\hat{z}^T A^T \hat{y} = \hat{s}^T M^T \hat{T}^{-1} t = u^T t \tag{28}$$

and that

$$\hat{z}^T \hat{z} = \hat{s}^T \hat{D} \hat{s} = v^T C \hat{s} = u^T \hat{T} u. \tag{29}$$

Eq. (28) is the sum of all measured traveltimes, while (29) is the sum of all traveltimes associated with the current best estimates of the slowness $\hat{s}$ and the ray-path matrix M. Thus, the condition (26), or equivalently (23) with $\gamma = 1$, implies that

$$u^T t = u^T \hat{T} u, \tag{30}$$

i.e., the hyperplane of (27) is the unique one for which the total measured traveltime is the same as the total traveltime for the trial model $\hat{s}$. Our conclusion is that both of the weighted least-squares solutions z_μ and $\hat{z}$ lie in the hyperplane determined by the limiting equality of the relaxed constraint (6). Furthermore, this statement is true for any positive value of μ.

The result (30) provides both a simple physical interpretation (constant total traveltime) and a simple geometrical interpretation (orthogonal to $\hat{z}$) for the hyperplane containing both z_μ and $\hat{z}$. In fact, it is this result that motivated us to pick this particular weighted least-squares problem (24) and the eigenvalue problem (20). We will term this choice of weights the "natural" one for the nonlinear traveltime tomography problem.

4. Modified ridge regression algorithm

Figure 4 illustrates the ideas presented in sections 2 and 3 and will also help to clarify the ideas underlying the new algorithm to be developed in this section.

The key ideas behind our new algorithm may now be summarized as follows: Given a set of transmitter-receiver pairs and any model slowness s, Fermat's principle may be used to find the ray-path matrix M associated both with s and with any slowness γs (where $\gamma > 0$) in the same direction as s. An optimum scale factor γ may be found by doing a weighted least-squares fit to the traveltime data. Good weights to use for low contrast reconstructions are described in detail in Berryman [20]. For high contrast applications, good weights were presented in section 3.3.

Having found the optimum slowness $\hat{s} = \gamma s$ in the given direction, we next attempt to improve the model by finding another direction in the slowness vector space that gives a still better fit to the traveltime data. As many others have done, we first compute a damped least-squares solution $s_\mu = \hat{D}^{-\frac{1}{2}}\hat{z}_\mu$ using (25). Next we note that both of the points found so far are guaranteed to lie in the nonfeasible part of the vector space – at least one and generally about half of the ray paths for both of these models will have traveltimes shorter than that of the measured data. Furthermore, although the point s_μ gives a better fit to the traveltime data, this fit is certainly spurious to some extent because it is based on the wrong ray-path matrix; the ray-path matrix M used in the computation of s_μ from $\hat{s}$ is the one that was correct for slownesses along the direction $\hat{s}$. Thus, both of the points we have found so far lie on the nonfeasible side of the feasibility boundary and the second point s_μ is of questionable worth because its value was also obtained in an essentially inconsistent manner.

We wish to stress this point: The motivation for performing a least-squares fit to the traveltime data is based on an implicit assumption of linearity, *i.e.*, closeness between the current model and the next model found after the correction step. If this linearity assumption is violated (as it often will be in the problems with high contrast we are considering), then our motivation for improving the least-squares fit to the data is not so strong and other criteria for choosing the size and direction of the correction step should be considered.

Now recall that the solution of (1), if one exists, must lie *on* the feasibility boundary. So we would like to use s_μ and $\hat{s}$ to help us find a point on this boundary that is optimum in the sense that it is as consistent as possible (i)

with the ray-path matrix M, (ii) with the measured traveltimes t, and (iii) with the feasibility constraints. The fact that traveltime error may be reduced by moving in the direction of s_μ may still give us an important clue about the best direction to move in the vector space, *i.e.*, we may want to move in the direction $s_\mu - \hat{s}$ but perhaps we should stop before arriving at s_μ. How far then should we move in this direction?

If we consider Figure 4, we are reminded that the feasible region is convex. Therefore, there may exist a point s_m between the points s_μ and $\hat{s}$ that is closer to the feasible region than either of the two end points. If we could find this point s_m and then scale up to the point in the same direction lying on the feasibility boundary, then we have found s_f in the figure. In principle, it is possible to find the point on the line $s_\mu - \hat{s}$ closest to the feasibility boundary. However, it is much easier to compute the feasibility violation number $N_M(s)$. As we move in the direction $s_\mu - \hat{s}$ from $\hat{s}$, we generally find that this number achieves a minimum value at some intermediate point. This point of minimum $N_M(s)$ is the point s_m in the figure.

Now we will prove that all three of the points $\hat{s}$, s_μ, and s_f are distinct unless we have found an exact solution to the inversion problem. Consider the possibility that $\hat{s}$ satisfies (1) so that

$$M\hat{s} = \hat{T}u = t. \tag{31}$$

Then, $M^T u = \hat{D}\hat{s} = M^T\hat{T}^{-1}t$ which is equivalent to $\hat{D}^{\frac{1}{2}}\hat{s} = (\hat{D}^{-\frac{1}{2}}M^T\hat{T}^{-\frac{1}{2}})\hat{T}^{-\frac{1}{2}}t$ or, in canonical form,

$$\hat{z} = A^T\hat{y}. \tag{32}$$

Eq. (32) implies that the right side of (25) vanishes. Since the matrix $(A^TA + \mu I)$ is nonsingular, (32) and (25) imply that

$$z_\mu = \hat{z}. \tag{33}$$

Thus, (31) implies (33). Similarly, working backwards through this line of reasoning, we see that (33) also implies (31). Our conclusion is that $z_\mu = \hat{z}$ if and only if $\hat{s}$ is an exact solution of (1). This same conclusion is reached by noting that the absolute minimum (*i.e.*, zero) of the objective function (24) is attained only when $s = \hat{s}$ and $Ms = t$. From previous arguments, we also know that the least-squares solutions z_μ and $\hat{z}$ never lie on the feasibility boundary unless the corresponding slownesses are exact solutions of (1). So unless we have already solved the problem, these three points form a triangle and the size of the triangle provides a measure of our distance from a solution.

5. Results

These ideas have all been repeatedly confirmed in a large number of reconstructions on synthetic examples. Figure 5 shows some examples of the ray paths used both in forward modeling and in the reconstructions. The model shown has two anomalies, a slow anomaly on top and a fast anomaly on the bottom. The percentage contrasts in these three examples are: 20%, 50%, and 100%. The model contains 8×16 square cells. The method used to generate the ray paths is new, although it is essentially a simplified version of an algorithm by Prothero *et al.* [23]. We represent the graph of a ray by a straight line plus a Fourier sine series expansion. The end points of the straight line are the locations of the transmitter and the receiver. The coefficients of the sine functions are found using a simplex search routine [24] to minimize the total traveltime along the ray path. For all examples shown in this paper, only the first two coefficients in the sine series are allowed to vary — thus severely limiting the possible shapes of the ray paths that can be found, but also limiting the amount of computation time required to obtain these approximations to the true ray paths.

The remaining Figures (6-8) show examples of reconstructions for three high contrast models having anomalies with 20%, 50%, and 100% contrasts. The traveltime data consists of 320 rays, including 256 (16×16) from left to right and 64 (8×8) from top to bottom. This measurement configuration was chosen to minimize the effects of ghosts, since the main issue being addressed in this paper is the possibility of obtaining reconstructions of high contrast (nonlinear) anomalies. The standard damped least-squares results are shown first (a) and the new results based on the algorithm using the minimum feasibility violation number to determine the correction step size are shown second (b). The layout of each frame is the same as Figure 5 with the slow anomaly at the top and the fast anomaly at the bottom. Reading from left to right and top to bottom the first five frames are the reconstructed slowness models at iterations 1, 11, 21, 31, and 41. The sixth frame (lower right hand corner) is the ideal (target) slowness model.

The two methods produce very similar results for the 20% anomalies shown in Figure 6. However, the new method produces a very stable accurate reconstruction while the damped least-squares method starts to diverge somewhat in the later stages — overshooting in the slow region and producing some vertical stripes in the homogeneous region that are no doubt caused by contributions from eigenvectors with low eigenvalues. The results start to diverge significantly for the 50% anomalies in Figure 7. The damped least-squares results are very noisy already by iteration 11 and are clearly not very useful in this form after more iterations. (A smoothed version of these reconstructions will at least show the locations of the slow and fast anomalies quite clearly.) The best reconstruction using the damped least-squares method is actually obtained after just a few iterations [15], so this method should be terminated quickly. However, for real data we do not know when to terminate the iteration process, so the damped least-squares method is not robust. By contrast, the new method produces results for 50% contrast that are quite reasonable: the fast anomaly is reconstructed very well while the slow anomaly is located well and general features are reproduced. The slow anomaly is always harder to reconstruct for high contrasts because few or no first arrivals pass through this region. Thus, first arrival data covers this area poorly. The results obtained by the two methods are startling different in Figure 8 (100% contrast). The damped least-squares method is quite unstable for such high contrasts. In fact, the method produces singular results by iteration 20, and the process terminates. The new method produces a stable result that retains the main features of the target model.

Our method converges quite rapidly to a definite result unless we force the algorithm to make a minimum percentage correction step (say 1-10% of the distance along $z_\mu - \hat{z}$) at each iteration. The triangle size decreases monotonically to a small value unless we force such corrections to be made; then the triangle size can decrease monotonically at first and subsequently oscillate around a small number.

The traveltime data used in these reconstructions were generated using the same ray tracing algorithm that is used in the reconstruction itself, *i.e.*, our simplex search algorithm. The results of this algorithm have been compared to the results obtained using Vidale's new finite difference approach to first arrival traveltime computation [25]. For small contrasts, both methods produce uniformly excellent results ($< 1\%$ error) with comparable short computation time. For larger contrasts, both algorithms become less accurate: (1) The simplex ray path is constrained by the use of only two coefficients in the sine series expansion to be quite smooth, perhaps smoother than it should be for such high contrast media. (2) The finite difference computation does not allow for all possible spreading patterns of wave energy, and therefore tends to miss the fastest ray paths in regions of high contrast where multiply reflected waves and other wave trapping effects can produce unusual results. We have tried using the results of Vidale's traveltime computations as the data for our reconstructions. For the cases of small contrast anomalies, the results are virtually the same. For higher contrast anomalies, Vidale's traveltimes used as input produce reconstructions that are somewhat noisier looking than those presented here, but basically the results are comparable. We also tried a hybrid method of finding the traveltime data: After generating the traveltime data using both the simplex method and Vidale's method, we construct a hybrid data set that contains the smaller of these two traveltimes between each transmitter and receiver. When used in the reconstructions, these hybrid data produced reconstructions that are again comparable to those shown here, but they are somewhat less noisy than these results — indicating that the traveltime data are indeed better in this case.

These ideas have also been tested on real seismic and real electromagnetic data. In all cases tested, the method converges to a reasonable model of the wave speed structure. Detailed analysis for applications of these methods to real data will be presented elsewhere.

6. Discussion and conclusions

For real problems with high contrasts and noisy data, the damped least-squares method does *not* converge and we never know when to terminate the iteration sequence. By contrast, our new algorithm converges quickly (in 15-20 iterations) to a solution in its convergence set (*i.e.*, not a single point, but a region of the model space with very similar characteristics). Stable iteration to such a convergence set is the most that could be expected when the traveltime data have errors and are therefore inconsistent.

The main purpose of this paper has been to develop a simple geometrical picture of the model space relevant to the iterative reconstruction algorithms. We have shown that Fermat's principle leads to one stable algorithm — a modified ridge regression method that is quite successful at reconstructing high contrast wave speed anomalies. This one algorithm certainly does not exhaust the possibilities. Other algorithms could also be envisioned based on the picture we now have of the model slowness vector space (Figure 3). With the advent of massive parallel processing computers, we should be able to map the feasibility surface and then use nonlinear programming techniques [26] or some sophisticated minimization technique such as simulated annealing [27] to produce still better reconstructions based on the general ideas presented here.

Acknowledgments

I thank J. A. Beatty and T. J. Yorkey for their contributions to earlier versions of the reconstruction codes. I thank J. Vidale for supplying me with a copy of his finite difference algorithm for computing first arrival traveltime data. Work performed under the auspices of the U. S. Department of Energy by the Lawrence Livermore National Laboratory under contract No. W-7405-ENG-48 and supported specifically by the LLNL Engineering Research Program, by the LLNL Earth Sciences Department Institutional Research and Development Program, and by the Geosciences Research Program directed by the DOE Office of Energy Research within the Office of Basic Energy Sciences, Division of Engineering and Geosciences.

References

[1] Herman G T 1980 *Image Reconstruction from Projections — The Fundamentals of Computerized Tomography* (New York: Academic) pp 1-18, 103-106

[2] Natterer F 1986 *The Mathematics of Computerized Tomography* (New York: Wiley) pp 1-8, 102-157

[3] Hounsfield G N 1973 *Brit. J. Radiol* **46** 1016; Cormack A M 1963 *J. Appl. Phys.* **34** 2722

[4] Gordon R, Bender R, and Herman G T, 1970 *J. Theor. Biol.* **29** 471

[5] Schomberg H 1978 *J. Phys. D: Appl. Phys.* **11** L181; Kak A C 1984 in *Digital Image Processing Techniques* (New York: Academic) edited by Ekstrom M P Chapt. 4 pp 111-170

[6] Bracewell R N and Riddle A C 1967 *Astrophys. J.* **150** 427

[7] Aki K, Christoffersson A, and Husebye E S 1976 *J. Geophys. Res.* **82** 277

[8] Aki K and Richards P G 1980 *Quantitative Seismology — Theory and Methods* Vol II (San Francisco: Freeman) Sec. 12.3

[9] Lager D L and Lytle R J 1977 *Radio Sci.* **12** 249

[10] Bois P, LaPorte M, Lavergne M, and Thomas G 1972 *Geophysics* **37** 471

[11] Dines K A and Lytle R J 1979 *Proc. IEEE* **67** 1065; Lytle R J and Dines

K A 1980 *IEEE Trans. Geosci. Remote Sensing* **GE-18** 234

[12] Justice J H, Vassiliou A A, Singh S, Logel J D, Hansen P A, Hall B R, Hutt P R, and Solanki J J 1989 *Geophysics: The Leading Edge of Exploration* **8** 12

[13] Wolf E 1969 *Optics Commun.* **1** 153; Mueller R K, Kaveh M, and Wade G 1979 sl Proc. IEEE **67** 567; Devaney A J 1984 *IEEE Trans. Geosci. Remote Sensing* **GE-22** 3

[14] Born M and Wolf E 1959 *Principles of Optics* (London: Pergamon) pp *xix-xx*, 127-129

[15] Berryman J G 1989 *Phys. Rev. Lett.* **62** 2953

[16] Grünbaum F A 1980 *Numer. Funct. Anal. Optimiz.* **2** 31

[17] Louis A K 1981 *Math. Meth. Appl. Sci.* **3** 1

[18] Strang G 1986 *Introduction to Applied Mathematics* (Wellesley, Mass: Wellesley-Cambridge Press) pp 673-683

[19] Lanczos C 1961 *Linear Differential Operators* (New York: Van Nostrand) Chapt 3 pp 100-162

[20] Berryman J G 1989 *IEEE Trans. Geosci. Remote Sensing* **27** 302

[21] Marquardt D W 1963 *J. Soc. Indust. Appl. Math.* **11** 431; Marquardt D W 1970 *Technometrics* **12** 591

[22] Titterington D M 1985 *Intern. Stat. Rev.* **53** 141

[23] Prothero W A, Taylor W J, and Eickemeyer J A 1988 *Bull. Seism. Soc. Am.* **78** 1190

[24] Nelder J A and Mead R 1965 *Computer J.* **7** 308

[25] Vidale J 1988 *Bull. Seismol. Soc. Am.* **78** 2062

[26] Fiacco A V and McCormick G P 1968 *Nonlinear Programming: Sequential Unconstrained Minimization Techniques* (New York: Wiley) Ch 6

[27] Kirkpatrick S, Gelatt C D Jr, and Vecchi M P 1983 *Science* **220** 671

FIGURE CAPTIONS

Fig. 1. Schematic diagram illustrating an experimental configuration typical of cross borehole seismic or electromagnetic tomography. For linear inversion, the straight rays shown are the ones normally used for backprojection.

Fig. 2. Diagram to illustrate terms in Eq. (1). Rectangular cells are numbered from 1 to n, with constant slownesses $s_1, \ldots, s_n$. The ray length for path i through cell j is $l_{i,j}$. The total traveltime for path i is $\sum l_{i,j} s_j = t_i$.

Fig. 3. Illustration of the distinction between local and global feasibility boundaries. Local feasibility is determined by the traveltime data vector and one ray-path matrix M. Global feasibility is determined by t and all possible Ms. Both types of feasibility region are convex.

Fig. 4. Illustrating the main points of the new reconstruction algorithm. The variable z_j is the wave slowness s_j weighted by the square root of the modified cell coverage $\hat{D}_j$ [see Eq. (20)]. The axes are the weighted slownesses for any two cells (j and k) in the model space. Point $\hat{z}$ is the initial value for the next step of the iteration scheme. Point z_μ is the damped weighted least-squares solution. Point z_m is a linear combination of $\hat{z}$ and z_μ chosen because it has the smallest number of feasibility violations. Point z_f is the unique point on the feasibility boundary in the same direction as z_m.

Fig. 5. Three examples of high contrast slowness models to be reconstructed from synthetic traveltime data. In each frame the top anomaly is slower than the background while the bottom anomaly is faster than the background. Percentage contrast from left to right: (a) 20% , (b) 50%, and (c) 100%. Superimposed on the models are examples of the curved ray paths obtained using the simplex search method and used in the reconstructions.

Fig. 6. Slowness reconstructions of 20% anomalies at iteration numbers 1, 11, 21, 31, and 41 for two methods: (a) the standard damped least-squares method and (b) the new method using minimum feasibility violation number to determine correction step size. The ideal (target) slowness model is shown in the lower right frame.

Fig. 7. As in Fig. 6 for 50% anomalies.

Fig. 8. As in Fig. 6 for 100% anomalies, except that the iteration frames shown in (a) are 1, 6, 11, 16, and 19.