

Application of the POCS inversion method to cross-borehole imaging

Ingi Th. Bjarnason* and William Menke‡

ABSTRACT

Cross-borehole tomography suffers from a well-known problem of data incompleteness: the limited ray coverage dictated by the poor experimental geometry implies that certain features of the velocity field are not determined by the data. Construction of a tomographic image of the velocity field therefore requires the addition of prior constraints to the inversion. In the Fourier wavenumber domain (assuming straight-line rays), the process of adding prior constraints is equivalent to specifying unmeasured wavenumber coefficients. The projection onto convex sets (POCS) algorithm can impose physically plausible constraints that allow high quality tomographic images to be produced. Each constraint is viewed as defining a set (in function space) of images that satisfy that particular constraint. The POCS method finds one or more images in the intersection of the constraining sets, which is equivalent to finding an image that simultaneously satisfies a number of constraints in-

cluding the observed data. The sets of images that we employ include: those that satisfy the data in the sense of having certain known wavenumber components, those that have bounded energy in certain unmeasured wavenumber components, those that have seismic velocity bounded everywhere (e.g., non-negative), and those in which the velocity structure is confined to the region between the boreholes. An advantage of the POCS algorithm is that it allows both space-domain and wavenumber-domain constraints to be imposed simultaneously. In our implementation of the POCS algorithm, we make use of the fast Fourier transform to rapidly iterate between the space and Fourier-wavenumber domains. We test the method on synthetic data, and show that it significantly reduces the artifacts in the image, when compared to other methods. We then apply it to data from a cross-borehole experiment in Manitoba, Canada, that were previously studied by others. We achieve a tomographic image of the velocity field that is similar in many respects to the results of others, but which possesses fewer artifacts.

INTRODUCTION

Images of physical parameters of the earth almost always suffer from distortions introduced by incomplete data coverage. The incompleteness is usually due to difficulties in achieving an experimental design that completely samples the earth's interior, such as in borehole tomography, where it is difficult to measure near-vertical rays (see, for instance, Menke, 1984). Incompleteness can also result from inherent insensitivity of the data to the model, such as in gravity surveys, where it is well known that the earth's density distribution cannot be recovered even when complete, noise-free gravity measurements are available.

Several different approaches are available for dealing with incompleteness when producing tomographic images of the earth. Backus and Gilbert (1970), for instance, recommend that the earth always be viewed through a filter that blocks out unresolved features in the image. Vasco (1986) recommends that the extremal bounds of physical parameters be imaged, rather than the physical parameters themselves. Yet the most widely applied approach—and the one examined here—is to specify prior constraints on the properties of the imaged parameters. Of critical importance is the type of prior information being imposed and its physical reasonableness in the context of the particular problem.

Manuscript received by the Editor December 4, 1991; revised manuscript received October 8, 1992.

*Formerly Lamont-Doherty Earth Observatory and Dept. of Geological Sciences, Columbia University, Palisades, N.Y. 10964; presently Carnegie Institute of Washington, Dept. DTM, 5241 Broad Branch Rd., N. W., Washington, DC 20015.

‡Lamont-Doherty Earth Observatory and Department of Geological Sciences, Columbia University, Palisades, NY 10964.

© 1993 Society of Exploration Geophysicists. All rights reserved.

Many tomographic inversion techniques supplement the data with some kind of prior constraint. The well-known damped least-squares method (Jackson, 1979; Tarantola and Valette, 1982; Menke, 1989) is one such technique. It adds a minimum energy or smoothness constraint. Other more specialized data extrapolation techniques also add prior information, as in the bandwidth-extension technique of Gerchberg (1974) and Papoulis (1975), and the image recovery method of Youla and Webb (1982).

The underlying mathematics of these methods (and many others) can be described as projection onto convex sets (POCS) (Halperin, 1962; Bregman, 1965; Gubin et al., 1967; Menke, 1991). The image of the earth, which is unknown, is viewed as a member of an infinite set of possible images. The data form one constraint that helps to select this image from the set of all possible images—it must be within the subset of images that satisfy the data. Each prior constraint corresponds to a different set. For instance, if the image is constrained to be nonnegative everywhere, then it must be in the set of nonnegative images. The intersection of all the sets defined by the data and the prior constraints define the set of possible solutions to the tomographic problem. Furthermore, as more and more constraints are added, the region of intersection becomes smaller and smaller and the set of acceptable solutions may shrink toward a unique solution. POCS provides an elegant algorithm for finding the intersection of overlapping sets.

Cross-borehole tomography is a widely used procedure to image the seismic velocity (or slowness) in the plane between two boreholes (Mason, 1981; McMechan, 1983; Wong et al., 1983; Menke, 1984; Bregman et al., 1989). As these and other authors have noted, the experimental geometry of the cross-borehole experiment lacks near-vertical incidence rays, so horizontal resolution suffers. On the other hand, many different kinds of prior constraints, such as the typical ranges of seismic velocities for crustal rocks and the typical magnitudes of velocity heterogeneities are available for cross-borehole experiments.

EXPOSITION OF THE PROBLEM

The cross-borehole tomography problem is to image the acoustic slowness $s(x, y)$ in the plane between two boreholes based on measurements of the traveltime T made between the boreholes. If we accept that the rays connecting the source and receiver are straight lines, which is appropriate for weakly heterogeneous structures, then the traveltimes are related to the acoustic slowness by the well-known Radon transform:

$$T(u, \phi) = \int_{-\infty}^{+\infty} dx \int_{-\infty}^{+\infty} dy s(x, y) \times \delta(s - x \cos \phi - y \sin \phi). \quad (1)$$

The independent variables u and ϕ index each ray by its midpoint position and orientation, respectively. It has long been known that $s(x, y)$ can be uniquely determined from $T(u, \phi)$ if T is measured at every possible u and ϕ (Radon, 1917). In practice, cross-borehole experiments can only measure traveltimes for subhorizontal rays, so T is typically known for a much smaller range of orientations than $0 \leq \phi < \pi$. Hence

images obtained from cross-borehole experiments are inherently nonunique.

The traveltimes data corresponding to rays all having the same orientation (value of ϕ) are conventionally called a projection. If we assume (as we will in the rest of this paper) that only complete projections are unmeasured, then the nature of the nonuniqueness can be clarified by examining the Fourier slice theorem. This theorem relates the two-dimensional (2-D) Fourier transform of the slowness, $s(k_x, k_y)$, to the one-dimensional (1-D) Fourier transform of the Radon transform, $T(k_u, \phi)$:

$$T(k_u, \phi) = s(k_x \rightarrow k_u \cos \phi, k_y \rightarrow k_u \sin \phi). \quad (2)$$

Here, k_x , k_y , and k_u are spatial wavenumbers. When $T(u, \phi)$ is known for only a limited number of projections, the Fourier transform of the image is known only in a fan-shaped portion of the image's wavenumber plane. Wavenumber components outside the fan are completely undetermined.

One approach to producing an acceptable image is to specify some prior constraints that have the effect of "filling in" the undetermined wavenumber components of the image. Note, however, that simply setting these values to zero is unacceptable, since the sharp edges in the wavenumber plane yield an image with pronounced streaks [for the same reasons that filtering a (1-D) time series with a narrow boxcar leads to very "ringy" results].

Failure of minimum norm methods

One class of prior constraint that is very widely used in inverse theory is that the unknown wavenumber components be chosen to minimize some measure of the complexity of the image, as measured by some appropriate L_2 norm:

$$\min \int_{-\infty}^{+\infty} dx \int_{-\infty}^{+\infty} dy |Ls(x, y)|^2. \quad (3)$$

The linear differential operator L can be chosen in a variety of ways. For instance, the choice $L = d^2/dx^2 + d^2/dy^2$ selects the image with the least curvature. Unfortunately, when the data are incomplete because certain projections have been omitted, this minimization principle selects the unacceptable zero solution regardless of the choice of L . This result follows from Parseval's theorem:

$$\int_{-\infty}^{+\infty} dx \int_{-\infty}^{+\infty} dy |Lf(x, y)|^2 \propto \int dk_x \int_{R_1} dk_y |p(k_x, k_y)s(k_x, k_y)|^2 + \int dk_x \int_{R_2} dk_y |p(k_x, k_y)s(k_x, k_y)|^2 \quad (4)$$

Here $p(k_x, k_y)$ is the Fourier transform of L , and we have divided the wavenumber plane into two regions, R_1 and R_2 , inside and outside the fan, respectively. The R_1 integrand is specified exactly by the data, while the R_2 integrand is completely unknown. Since both integrands are positive everywhere, the minimum occurs when the second is set to

zero. Hence the left-hand side in equation (4) is minimized when $s(k_x, k_y) = 0$ inside R_2 .

Prior constraints implemented with POCS

POCS provides a method of imposing certain prior constraints, including nonlinear ones (see Gubin et al., 1967). The method views each constraint as defining a set (in function space) of images—the set of images that satisfies that constraint. The data themselves define one set: the set of images that satisfy the available traveltime data. Various other prior constraints can be, as we will delineate below, associated with other sets. The solution to the imaging problem is then equivalent to finding the one or more images that is in the area of intersection of all the sets. The solution is unique if the intersection contains only one image. Even one example of a nonunique solution may be of use, if the constraints force the many possible solutions to differ only in inconsequential details. The application of POCS to problems in geophysical inversion has been discussed by Menke (1991).

POCS provides an iterative method of finding one point (in function space) in the region of intersection of the sets—that is, one image that satisfies all of the constraints. The only requirement is that each of the sets be “convex,” meaning that if two images, s_1 and s_2 , are within a set then their linear combination, $\alpha s_1 + (1 - \alpha)s_2$ is in it too, for $0 < \alpha < 1$. Many interesting sets have this property. We begin with an initial guess for the image, not necessarily within any of the sets, and iteratively project it orthogonally upon the surfaces of the sets. Provided all the sets are convex, the projected image eventually converges to the region of the intersection of the sets. The sets that we will apply in this paper are listed below.

Set K (for wavenumber): Since the data are assumed to be in the form of complete projections, we can represent them either in the Radon (i.e., projections) or Fourier domain (i.e., wavenumber coefficients). We choose the later, since it allows us to use the fast Fourier transform (FFT) algorithm during the iterations. The data therefore define set K of all images with a prescribed Fourier transform $s(k_x, k_y)$ inside a fan shaped region R_1 . The projection algorithm for this set is very simple: the image is transformed into the wavenumber domain using the FFT, those values within the fan are reset to their prescribed values, and the image is fast Fourier transformed back.

Set C (for compact support): A significant difference between the way traveltime is computed in the Radon transform and in the actual experiment is that the rays in the former are of infinite length, while in the latter they are bounded by the source and receiver points. This does not present any difficulty—the two are equivalent as long as we define the image to be zero outside of the region between the boreholes. Inversion of the Radon transform of complete, noise-free data will indeed recover a zero image everywhere outside the region between the boreholes. On the other hand, real data is incomplete and noisy, so the finiteness of the rays should be included as a prior constraint by forcing the image to have “compact support,” that is, forcing it to be zero outside the region between the boreholes. The image

is projected onto this set simply by setting the prescribed pixels to zero.

Set B (for bounded): Since the acoustic slowness is an inherently positive quantity, we must choose $s(x, y)$ from the set of images that are nonnegative everywhere. In practice, stronger bounds can be placed on the slowness, say $s_\ell < s(x, y) < s_u$, since lithological information about the rocks is often available. An image is projected onto this set by examining each pixel and resetting its value to the upper or lower limit should it fall below or above them, respectively. In this work we chose the allowable range of slowness perturbations to be twice the range of perturbations obtained by a standard, minimum-norm inversion of the data.

Set E (for energy): The total energy in the image can also be constrained to be less than a prescribed amount, that is:

$$\int_{-x}^{+x} dx \int_{-y}^{+y} dy |s(x, y)|^2 < E_0. \quad (5)$$

The image is projected onto this set by renormalizing its energy to E_0 if it exceeds this bound. In choosing the value for E_0 we made the assumption that the average energy of the wavenumbers in R_2 could be no greater than the average in R_1 .

Sets C and K alone are sufficient to specify a unique solution, since these two sets are guaranteed to intersect at no more than a single point in function space. (This result can be proven by applying the Paley-Wiener and Identity theorems of functional analysis). The result can be understood in the context of functional analysis in the following way: if the Fourier transform of a function with compact support is known in a patch, it can be analytically continued to all of the wavenumber plane. Hence sets B and E seem unnecessary. Their importance is seen when the function is discretized into pixels, as it must be for implementation on a digital computer. Then the constraint of compact support is applied only to a finite number of points (pixels) in the (x, y) plane and the constraint of prescribed wavenumbers is also only applied at a finite number of points in the (k_x, k_y) plane. Then, if the total number of constraints is less than the total number of pixels, the intersection of sets C and K may have an infinite number of elements. Then sets B and E serve to physically select reasonable elements, and hence to stabilize the convergence. In our calculations, the number of constraints C and K are 93 percent of the number of pixels.

One could of course, choose the number of C and K constraints to exactly equal, or exceed, the number of pixels. But then the region of intersection may be very sensitive to noise, or not exist at all. In this case, POCS will not converge to a unique solution. POCS can, however, be used to select a solution that is consistent with the observed wavenumbers (by making the last projection onto K), and that has approximately compact support. Again sets B and E serve to physically select reasonable solution.

Since POCS is a very general algorithm applicable to any collection of convex sets, it should be no surprise that many existing iterative algorithms drawn from applied mathematics can be reinterpreted as special cases of POCS. Such a reinterpretation is useful, since it lends insight into the nature of the algorithm and provides a simple way to prove their convergence: Since POCS, in general, converges

(Gubin et al., 1967), these more specialized algorithms also converge. The bandwidth-extension technique of Gerchberg (1974) and Papoulis (1975) is one example of a specialized form of POCS, as it is the widely used algebraic reconstruction technique (ART). In our opinion, however, tomographic inversion benefits from explicit recognition of the role of POCS, since how many physically important constraints can be implemented then become apparent.

NUMERICAL IMPLEMENTATION

One of the attractive features of the POCS algorithm is that it is very simple to implement numerically. The image is represented in discretized form as a rectangular array of slowness values, or equivalently, as a rectangular array of Fourier-wavenumber coefficients. (The FFT algorithm rapidly switches back and forth between the two domains). The projection operation then simply consists of modifying some of the elements of the arrays. For instance, a space-domain image is projected onto set C of compact support by simply setting some of the elements outside the specified region to zero. The image is projected onto set B of bounded values by examining each of the pixels and by resetting the value to the value of the bound if it is outside of the bound. The image is projected onto set E of bounded energy by simply computing the total energy (by summing up the squared value of all the pixels) and by rescaling their values (multiplying them by an appropriate constant) if the energy exceeds the bound. Similarly, a wavenumber-domain image is projected onto set K by simply resetting some of its Fourier coefficients to prescribed values.

The tomographic inversion algorithm therefore consists of starting with an "initial guess" for the image (we use the minimum energy solution) and successively projecting it onto the sets. In our implementation, we first project onto the space-domain sets, then transform the image to the wavenumber-domain using the FFT, project onto the wavenumber-domain sets, and then transform back to the space-domain with an inverse FFT. This process constitutes one iteration. Given enough iterations, the POCS algorithm is guaranteed to eventually converge to the intersection of all the sets. In practice, we find 10–50 iterations to be sufficient. The most important constraints that the inverse solution should satisfy are the observed wavenumbers or traveltimes. Therefore we choose the last projection in the POCS iterations to be on the K set.

One practical advantage of this kind of iterative algorithm is that new convex sets can easily be added to the tomographic inversion code by simply adding a new projection subroutine. No modifications are required in parts of the code concerned with the other sets.

Synthetic and actual examples

We first test the method with a hypothetical slowness structure (Figure 1a) and form synthetic traveltimes data from it by summing over all rays with angles within -40 degrees to $+33$ degrees of horizontal (Figure 1b). Although traveltimes have been portrayed as zero (black) in this figure, they are in fact nonzero but unknown. An inversion based on arbitrarily setting these traveltimes to zero (or equivalently setting the unknown wavenumbers in R_2 to

zero) leads to a poor image (Figure 1c). Note that the image contains both pronounced streaks arising from the sharp edge along the boundary between R_1 and R_2 , and "bleeding" of the image into the outer portion (beyond the boreholes) which is in fact zero. After 30 POCS iterations onto the sets C, B, E, and K (in consecutive order), many of the artifacts in the image have been significantly reduced, and the correct support has been restored without adding any noticeable new distortions (Figure 1d). Most importantly, there is an improvement in the horizontal resolution.

We use an example of data, recorded from a cross-borehole experiment in metamorphic rock in Manitoba, Canada, that have previously been studied by Wong et al. (1983) and Bregman et al. (1989). Bregman et al.'s (1989) inversion is shown in Figure 2a, for reference. The boreholes are approximately 350 m deep, and are separated by 175 m (Figure 2a). The seismic imaging was confined to the depth range of 100–350 m. Bregman (1986) has verified that straight-line ray approximation is reasonable for this data set. We picked 560 compressional wave traveltimes from the digital seismograms but used only the best quality traveltimes for our inversion, or a total of 375 traveltimes. The traveltimes error of the good quality data was determined to be less than ± 0.15 ms, which is about 0.5 percent of the typical traveltimes between the boreholes. We have found the Radon transform domain a very useful format for displaying the traveltimes data (Figure 2b), since it organizes the data so that traveltimes from rays with similar orientations and positions are plotted at adjacent positions. The gradual increase in velocity with depth is readily apparent, as are departures from vertical stratification. Furthermore, individual rays with uncharacteristically small or large traveltimes (i.e., "outliers") are easy to identify. The plot also brings out the degree of incompleteness of the data set, since most of the Radon transform domain is void of data. The different number of data in Figures 1b and 2b should also be noted: 256 000 and 375 points, respectively.

Unlike synthetic data, the actual data does not quite constitute complete projections as are required by our use of set K. We avoid this problem by interpolating the data onto a rectangular region of the Radon transform domain using continuous curvature splines (Smith and Wessel, 1990). Alternatively, we could use set R of images with Radon transforms specified at known points. Projection onto R is straightforward [see Menke's (1991) discussion of his set D of functions that satisfy a general linear inverse problem], but is somewhat more cumbersome numerically.

A minimum energy inversion based on arbitrarily setting the unknown wavenumbers in R_2 to zero (Figure 2c) compares quite well with the inversion obtained by Bregman et al. (1989, see Figure 2a). The overall velocity increase with depth is similar, as are many of the lateral velocity perturbations. Both solutions have undesirable streaking affects along raypaths. There are some differences between the two solutions. This was as expected, because there are a number of differences between the two algorithms: (1) Bregman et al. (1989) used damped least-squares inversion that minimizes a linear combination of the prediction error and the energy of the solution. Our solution is, however, the absolute minimum energy solution. (2) Bregman et al. (1989) used about 550 traveltimes for their inversion. We used only the best

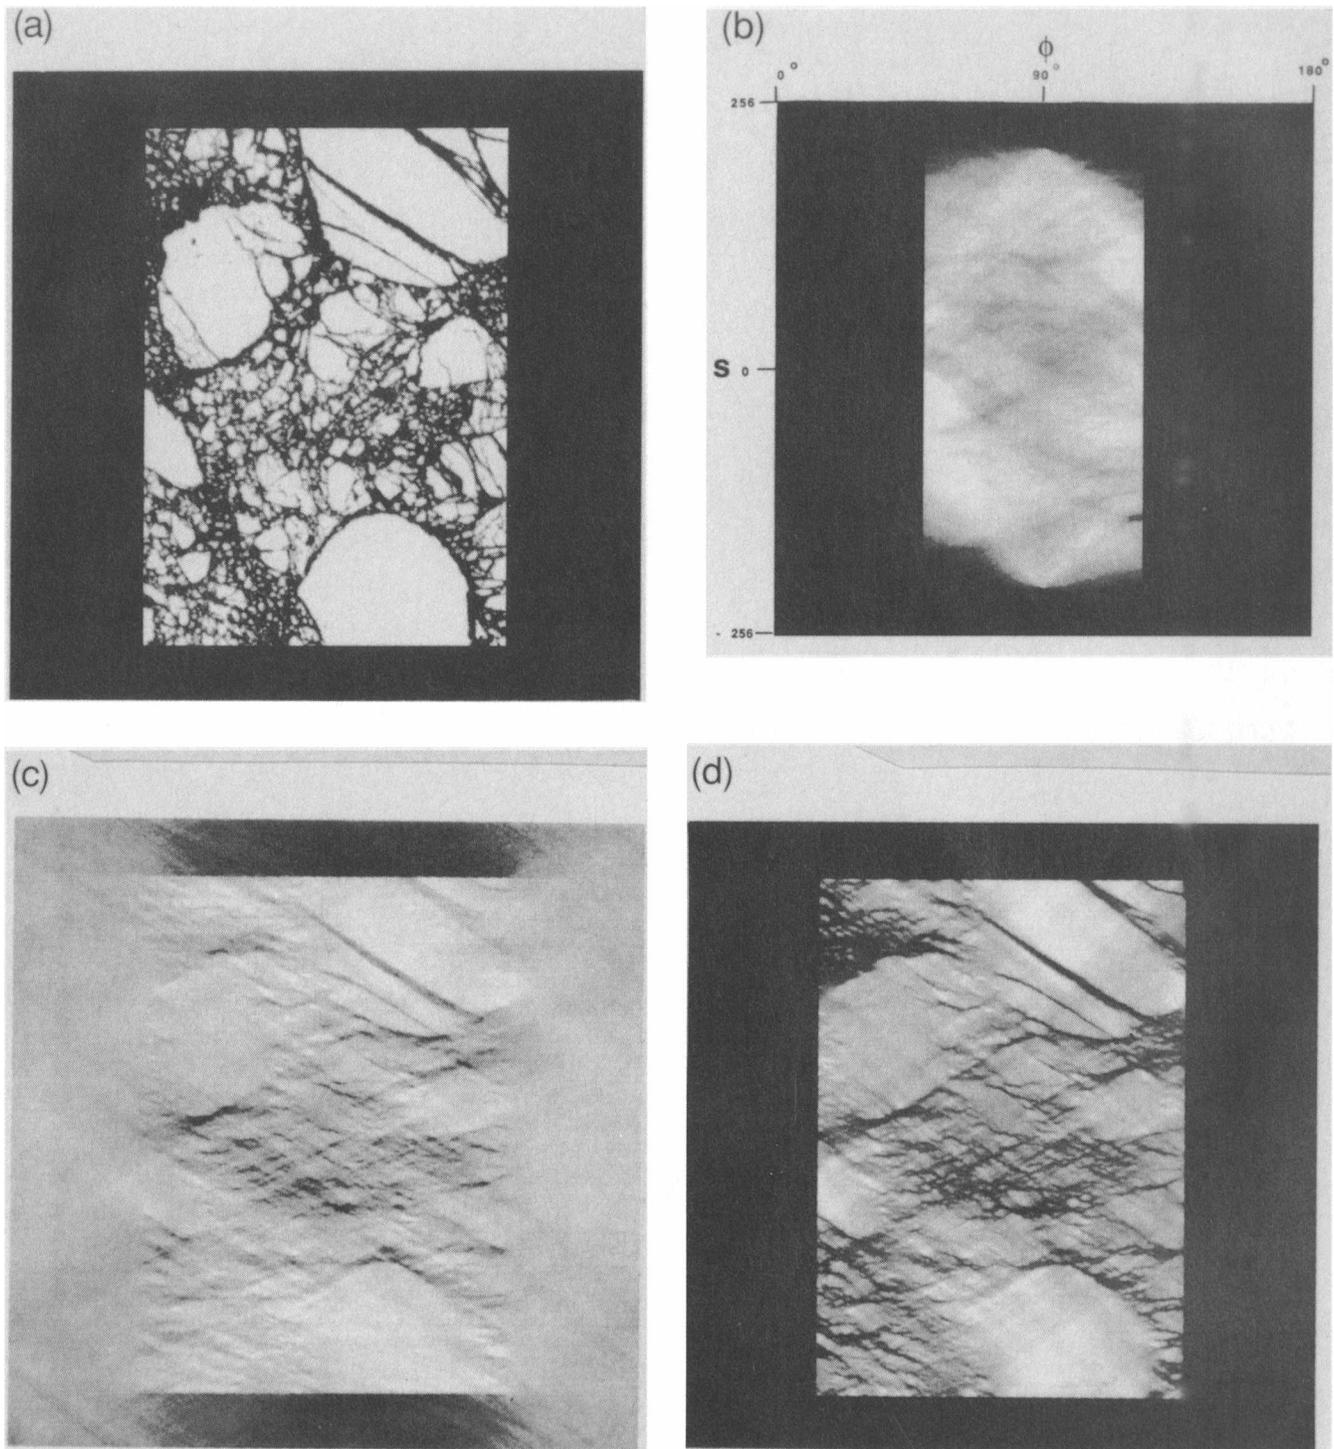


FIG. 1. Inversion of synthetic data. (a) A digitized photomicrograph of a fault gouge is used for the synthetic image. (b) Traveltime data for this image in the angular range -40 degrees to $+33$ degrees of horizontal. This range is similar to the range used in the Wong et al. (1938) cross-borehole experiment. Missing traveltimes are shown in black. (c) Reconstructed image of traveltime data using minimum energy inversion. (d) POCS inversion after 30 iterations. Note that the streaks present in the minimum energy inversion have been substantially reduced.

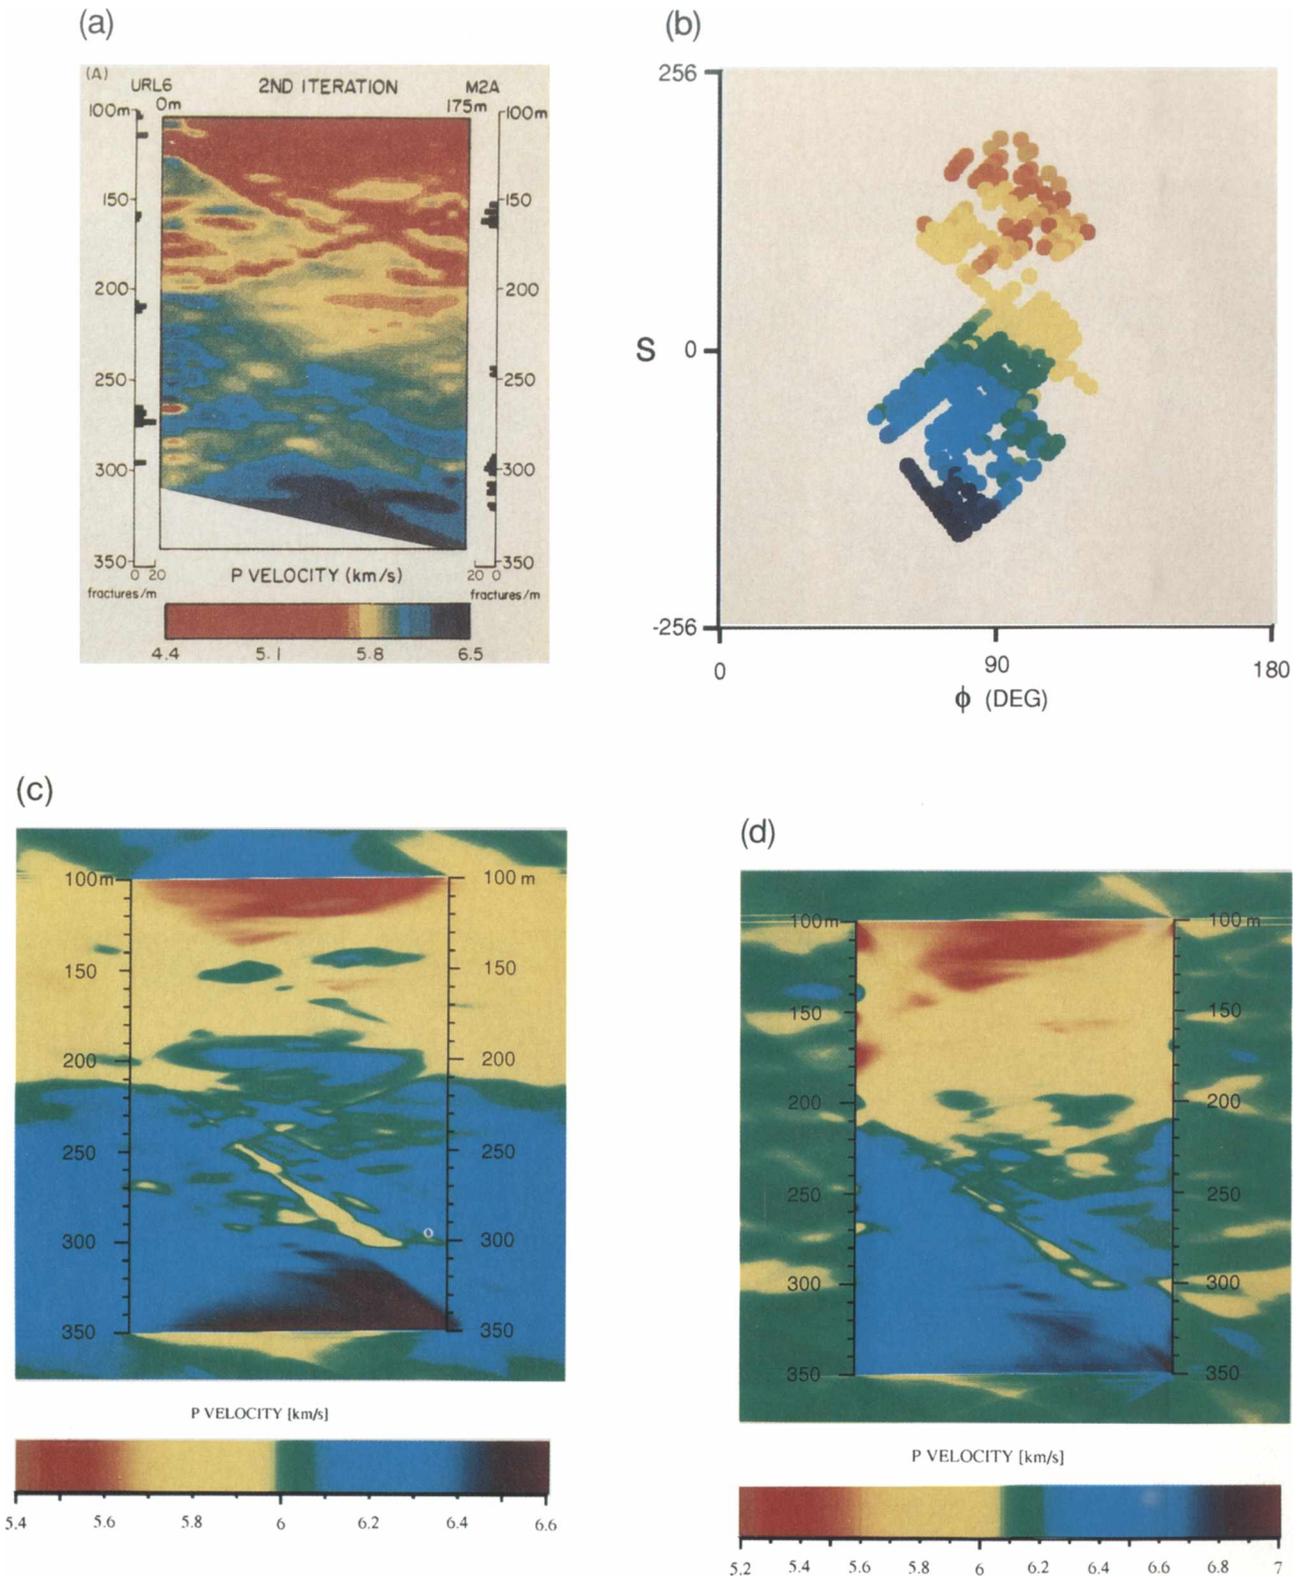


FIG. 2. Inversion of field data. (a) Experimental geometry showing boreholes. Fracture counts for the two boreholes are shown at the edges of the model. Bregman et al.'s (1989, Plate 1a) reconstructed structure is shown in the interior. (b) Traveltime data for this experiment. (c) Reconstructed image of traveltime data using minimum energy inversion. (d) POCS inversion after 10 iterations. Note the similarity with Bregman et al.'s (1989) inversion, except that the streaks present in the former have been substantially reduced.

quality data, or approximately 2/3 of the total data set. (3) Bregman et al.'s (1989) grid spacing is approximately one magnitude unit coarser than ours. (4) Bregman et al. (1989) ray trace through their perturbed solution model. We use a straight-line approximation of raypaths. We do not, however, think that this is an important factor in explaining the differences between the two solutions. Bregman (1986) showed that a straight ray approximation is reasonable for this data set.

The inversion after 10 POCS iteration onto sets C, B, E, and K (in consecutive order) is shown in Figure 2d. Compared to the minimum norm solutions (Figure 2a and 2c), the POCS inversion has, in general, reduced the streaks that are caused by the poor data coverage, and the solution seems to be smoother overall. Amplitude is reduced in most of the high-velocity perturbations at the top of the image, as well as in some of the low velocity perturbations in the lower half of the image. Consequently, the mean vertical compressional-wave velocity gradient increases somewhat, from 6.5 percent in Figure 2c to 9.5 percent in Figure 2d. There is also a slight increase in the size of the perturbations themselves, from ± 3.5 percent to ± 4.0 percent.

Some instabilities develop in the POCS inversion of these data. Suspicious oscillations of high and low perturbations develop along the left borehole (Figure 2d). The solution of Bregman et al. (1989) also suffers from this problem, but the oscillations of high and low perturbations along the left borehole are even more pronounced in their solution. Bregman (1986) ascribed them to the least-squares inversion placing high wavenumber oscillations at the least well constrained parts of the model to fit the data. Another instability develops in the POCS inversion where high-frequency streaks radiate into the image from the top corners of the boreholes. This is not observed in Bregman et al.'s (1989) inversion. The corner streaks in our solution have the same angles with respect to horizontal as the edges of the fan in the Fourier domain have with respect to vertical (33 and 40 degrees). Therefore, we believe that the source of the corner streaks are sharp discontinuities in high wavenumber values across the edge of the fan that have been erroneously amplified by the POCS inversion. However, streaks within the angular range of -40 degrees to $+33$ degrees of horizontal are caused by uncharacteristically large perturbations in the traveltimes of a single or only a few clustered rays compared to neighboring rays. Some of these streaks are probably caused by noise in the data. Others are, however, probably caused by the sparseness and incompleteness of the measurements. This has the effect of smearing large perturbations along raypaths over the image. It is important to note that POCS adds nonzero energy to the undetermined wavenumbers by trying to impose physically reasonable constraints, in contrast to the damped least-squares technique that adds nonzero energy to the undetermined wavenumbers without imposing any special physically reasonable constraints.

We ascribe the general instability problems of the POCS inversion of actual data to the effect of measurement noise. It is also possible that the wide interpolation that we performed in the Radon transform domain may have caused some "inconsistencies" in traveltimes, making the inversion less stable. In future work, we recommend projection onto

set R instead of K. In this way the interpolation in the Radon transform domain is achieved within the framework of POCS.

OTHER APPLICATIONS OF POCS IN TOMOGRAPHY

POCS can be used to apply a variety of other constraints to the inversion. For instance, the spectrum can be forced to be less than a model function, $|s(k_x, k_y)|^2 < S(k_x^2 + k_y^2)$. This constraint will yield a very smooth image when $S(k_x^2 + k_y^2)$ is taken to decay exponentially with radial wavenumber $k_x^2 + k_y^2$. Alternatively, very rough (fractal) images can be produced by choosing S to decay algebraically with wavenumber. This spectral constraint may be of particular importance if well-logs or seismic measurements from cores can be used to place statistical constraints on the roughness (high wavenumber spectrum) of the velocity between the boreholes.

The POCS algorithm can also be applied to a general tomographic inversion that employs curved rays, except that set D of functions that satisfy a general linear inverse problem (see above) must be used in place of K. While this change precludes the use of the computationally efficient FFT algorithm, it is still practical and may find occasional application.

CONCLUSIONS

The main objective of this work is to find useful methods to improve seismic tomographic images that are constructed from limited data. The POCS method permits a wide variety of prior information to be used to artificially increase data coverage. The only requirement is that the constraint form a convex set in model space. Physically important constraints can be effectively applied to the limited-coverage cross-borehole experiment. Particularly important are the constraints that the imaged seismic velocities fall within a prescribed physically reasonable range, and that the total energy of the velocity heterogeneities is bounded. Application of POCS iterative inversion to a synthetic data set with ray coverage similar to an actual tomography experiment led to very substantial reductions of artifacts in the tomographic image. The POCS method also reduces artifacts in actual data, but the measurement noise makes the inversion less stable.

Future research directions include the application of a wide range of physically plausible constraints, such as the spectral constraints mentioned above, and the implementation of an algorithm that is not dependent upon the use of the Radon transform and FFT and thus is applicable to tomography problems with curved rays.

ACKNOWLEDGMENTS

We thank Finnur Lárusson and Sigurthur Helgason for useful discussions of the underlying mathematics, Alberto Malinverno and David Caress for critical reviews of the manuscript, and Chris Marone for supplying us with the photomicrograph. We thank Atomic Energy of Canada and Peter Hurley for making the borehole data available to us. This work was supported by the National Science Foundation and the Office for Naval Research. This is Lamont-Doherty contribution number 5042.

REFERENCES

- Backus, G., and Gilbert, F., 1970, Uniqueness in the inversion of gross earth data: *Phil. Trans. Roy. Soc. London*, **A 266**, 123–192.
- Bregman, L. M., 1965, The method of successive projection for finding a common point of convex sets: *Soviet Math. Dokl.*, **6**, 688–692.
- Bergman, N. D., 1986, Tomographic inversion of crosshole seismic data: Ph.D. thesis, Univ. of Toronto.
- Bregman, N. D., Chapman, C. H., and Bailey, R. C., 1989, Traveltime and amplitude analysis in seismic tomography: *J. Geophys. Res.*, **94**, 7577–7587.
- Gerchberg, R. W., 1974, Super-resolution through error energy reduction: *Optica Acta*, **21**, 709–720.
- Gubin, L. G., Polyak, B. T., and Raik, E. V., 1967, The method of projections for the common point of convex sets: *USSR Comput. Math. and Math. Phys.*, **7**, 1–24.
- Halperin, I., 1962, The product of projection operators: *Acta Sci. Math.*, **23**, 96–99.
- Jackson, D. D., 1979, The use of a priori data to resolve nonuniqueness in linear inversion: *Geophys. J. Roy. Astr. Soc.*, **57**, 137–157.
- Mason, I. M., 1981, Algebraic reconstruction of a two-dimensional velocity inhomogeneity in the High Hazles Seam of Thoresby colliery: *Geophysics*, **46**, 298–308.
- McMechan, G. A., 1983, Seismic tomography of boreholes: *Geophys. J. Roy. Astr. Soc.*, **74**, 601–612.
- Menke, W. H., 1984, The resolving power of cross-borehole tomography: *Geophys. Res. Lett.*, **11**, 105–108.
- , 1989, *Geophysical data analysis: Discrete inverse theory*, 2nd ed.: Academic Press, Inc.
- , 1991, Applications of the POCS inversion method to interpolating topography and other geophysical fields: *Geophys. Res. Lett.*, **18**, 435–438.
- Papoulis, A., 1975, A new algorithm in spectral analysis and band-limited extrapolation: *IEEE Trans. Circ. and Sys.*, **CAS-22**, 735–742.
- Radon, J., 1917, Über die Bestimmung von Functionen durch ihre Integralwerte langs gewisser Mannigfaltigkeiten: *Berichte Sachsische Akademie der Wissenschafte, Leipzig, Math.-Phys. Kl.*, **69**, 262–277.
- Smith, W., and Wessel, P., 1990, Gridding with continuous curvature splines in tension: *Geophysics*, **55**, 293–305.
- Tarantola, A., and Valette, B., 1982, Generalized nonlinear inversion problems solved using the least-squares criterion: *Rev. Geophys. Space Phys.*, **20**, 219–232.
- Vasco, D. W., 1986, Extremal inversion of traveltime residuals: *Bull. Seis. Soc. Am.*, **76**, 1323–1345.
- Wong, J., Hurley, P., and West, G. F., 1983, Crosshole seismology and seismic imaging in crystalline rocks: *Geophys. Res. Lett.*, **10**, 686–689.
- Youla, D. C., and Webb, H., 1982, Image restoration by the method of convex projections: Part 1—Theory: *IEEE Trans. Medical Imaging*, **1**, 81–94.