

Automation of Shear-Wave Splitting Measurements using Cluster Analysis

by N. A. Teanby,* J.-M. Kendall, and M. van der Baan

Abstract The propagation of two orthogonally polarized shear waves, or shear-wave splitting, is arguably the most robust indication of seismic anisotropy in the Earth. This splitting can be parameterized in terms of the polarization of the fast shear-wave ϕ and the lag time between fast and slow components δt . These two parameters provide constraints on the mechanism causing the anisotropy. All methods of calculating splitting require a shear-wave analysis window to be selected. Then the ϕ and δt that best account for the splitting in that window are calculated. Conventionally the shear-wave analysis window is picked manually. However, manual window selection is laborious and also very subjective; in many cases different windows give very different results. We present a method for automating the selection of the window. First, the splitting analysis is performed for a range of window lengths. Then a cluster analysis is applied in order to find those measurements that are stable over many different windows. Once clusters of stable results have been found, the final choice of shear-wave analysis window corresponds to the measurement with the lowest error in the cluster with the lowest variance. Resulting estimates of ϕ and δt are objective, and very large datasets can be analyzed easily. The success of the technique is illustrated with application to a microseismic dataset of 324 events, which confirms previously published results using manually selected analysis windows.

Introduction

Seismic anisotropy is a ubiquitous feature of most geological materials. This is because rocks generally exhibit fabric or order, which leads to a directional dependence in seismic velocities. When a shear wave propagates through an anisotropic medium, energy is partitioned into orthogonally polarized fast and slow shear waves, which have similar waveforms. This property is referred to as “shear-wave splitting,” or sometimes “birefringence.” Splitting can be described in terms of two parameters: the fast polarization direction (ϕ) and the lag time between fast and slow shear waves (δt). These parameters provide information on the nature of the rocks that the wave propagated through. Anisotropy is caused by ordered features, such as fractures with preferential alignments, or lattice preferred orientation (LPO) of anisotropic minerals due to deposition or deformation. Therefore, measurements of anisotropy provide constraints on rock fabric and hence geological or geodynamical processes.

Shear-wave splitting is routinely measured in teleseismic data and is used to probe the fabric of crust and mantle rocks (Silver, 1996; Savage, 1999; Kendall, 2000). The

length scales of these observations vary from hundreds of kilometers for teleseismic *SKS* and direct *S* phases (Park and Levin, 2002; Wookey *et al.*, 2002) to kilometer scales for local shear waves (Gledhill and Stuart, 1996) and meter scales for microseismic events (Bokelmann and Harjes, 2000; Teanby *et al.*, 2004). Hence, splitting measurements can be used to interpret LPO, fracturing, and geodynamical processes over a large range of length scales.

A standard technique for measuring shear-wave splitting is the splitting correction method of Silver and Chan (1991) (see also Vinnik *et al.* [1989]). In this method, a shear-wave analysis window is selected manually and the ϕ and δt that best correct the splitting in this window are calculated by means of a grid search. The method will be described in more detail in the next section. A problem often encountered when using this approach is that the calculated ϕ and δt are sensitive to the choice of shear-wave analysis window. This means that manual selection of the shear-wave analysis window is subjective and can heavily influence the results. Automated window selection avoids this problem and gives an objective measurement.

Data volumes obtained in microseismic studies tend to grow quickly. In the study by Teanby *et al.* (2003), 324 located events were recorded on a single array of geophones over an 8-week recording time. Current deployments aim

*Present address: Atmospheric, Oceanic, and Planetary Physics, Department of Physics, University of Oxford, Clarendon Laboratory, Parks Road, Oxford, OX1 3PU, United Kingdom; teanby@atm.ox.ac.uk.

for multiple geophone arrays and longer recording times. Therefore, manual analysis of every event will be too time consuming. Hence, we require an automated and objective approach to still be able to handle such large data volumes in a convenient way. These large datasets should provide insights into lithological properties, making it possible to constrain fracturing and intrinsic anisotropy in geological models.

In this article, we attempt to automate the choice of analysis window by performing a grid search over different windows to find stable splitting measurements with small error. This should make the study of large datasets practical, while simultaneously reducing the subjectivity of measurements. We apply the method to the microseismic dataset from Teanby *et al.* (2003).

Automation of shear-wave splitting measurements has been previously attempted by Savage *et al.* (1989) and Aster *et al.* (1990). These methods have had some success with local earthquake data and illustrate the importance of automation. However, the disadvantage of these methods is that they do not address the effect that different shear-wave analysis windows can have on the results (see the Method section). Also, these early studies use methods that only use a small fraction of the waveform to find ϕ . Our technique has the advantage of investigating and trying to remove the dependence of results on the exact analysis window. We also use the full waveform technique of Silver and Chan (1991), making our results less susceptible to noise.

Method

We base our automated shear-wave splitting technique on the splitting correction method of Silver and Chan (1991). First, a shear-wave analysis window is defined, which is normally picked manually. If anisotropy is present, the particle motion within this window will be elliptical. Second, a grid search over ϕ and δt is performed, where both components are rotated by ϕ and one component is lagged by δt . The result that has the lowest second eigenvalue of the corrected particle-motion covariance matrix indicates linear particle motion after correction and is the solution that best corrects for the splitting. An F test is used to calculate the 95% confidence interval for the optimum values for ϕ and δt . Sandvol and Hearn (1994) suggested an alternative method for calculating errors in ϕ and δt based on a bootstrap analysis.

After the splitting correction has been applied, we require that the corrected waveforms in the analysis window match. The second eigenvalue of the particle-motion covariance matrix provides a measure of this match. The smaller the second eigenvalue, the better the match. We also require the errors on the splitting parameters to be as low as possible, and a threshold can be set in order to discard or keep results. Cycle skipping, caused by noisy or bandlimited data, results in multiple solutions with multiple values for ϕ and δt . A good result will have a unique solution. Criteria for reliable results were discussed in Savage (1999) and Silver and Chan

(1991). Figure 1 shows an example of a reliable splitting result.

Results are often sensitive to the exact choice of shear-wave analysis window. Before we automate window selection, it is important to consider what is a good choice of window. Most importantly, the window should be representative of the S wave and ideally be long enough to include several periods of the dominant frequency to prevent cycle skipping and decrease the influence of noise. However, it should not be so long as to include spurious secondary phases in later parts of the waveform, which will degrade splitting estimates. Results from windows smaller than one period tend to become unstable and have unrealistically low errors because only small fragments of a wavelength need to be matched. If there is no interference from other phases, the entire S -wave energy envelope can be used, giving excellent and very robust results (e.g., Fig. 1). Choosing the start of the window slightly before the onset of the shear wave often stabilizes results and reduces cycle-skipping effects, which are often a problem for signals with limited frequency content.

It is very important that the splitting parameters are stable over a wide range of different analysis windows. This stability ensures that the measurement is robust but is time consuming to test using manual window selection and is therefore not typically reported in shear-wave splitting studies. On the other hand, we use exactly this feature to select shear-wave analysis windows objectively and automatically. Figure 2 illustrates how slight changes in the shear-wave analysis window can drastically affect the splitting results.

To automate shear-wave splitting, we search over a range of window start/end times to find the optimum window for the splitting correction. Simply searching for the window that gives the lowest error bars in the measurements is not a good criterion as an unstable result, which is sensitive to small window changes, may be selected. Instead we search for measurements that are stable over many different analysis windows. This is achieved by varying the analysis window and looking for plateaus in δt and ϕ . These plateaus indicate stable splitting measurements, so that once the proper plateau has been identified the window that gives the smallest error can be selected. Figure 3 shows an example of splitting parameters obtained for 250 different shear-wave analysis windows. There is a stable plateau with small errors around window number 20, which indicates a stable solution. The problem is then to automate the selection of such plateaus.

Our method consists of three steps. First we calculate δt and ϕ for a range of start and end times of the shear-wave analysis window. In a 2D plot of ϕ against δt , stable regions of the window space will condense into points, or tight clusters, as shown in Figure 4. Second, we use cluster analysis to identify the stable regions. Because an unsupervised technique is required, the two main considerations are the type of clustering technique to use and a way to determine the optimum number of clusters. Finally, when a set of clusters

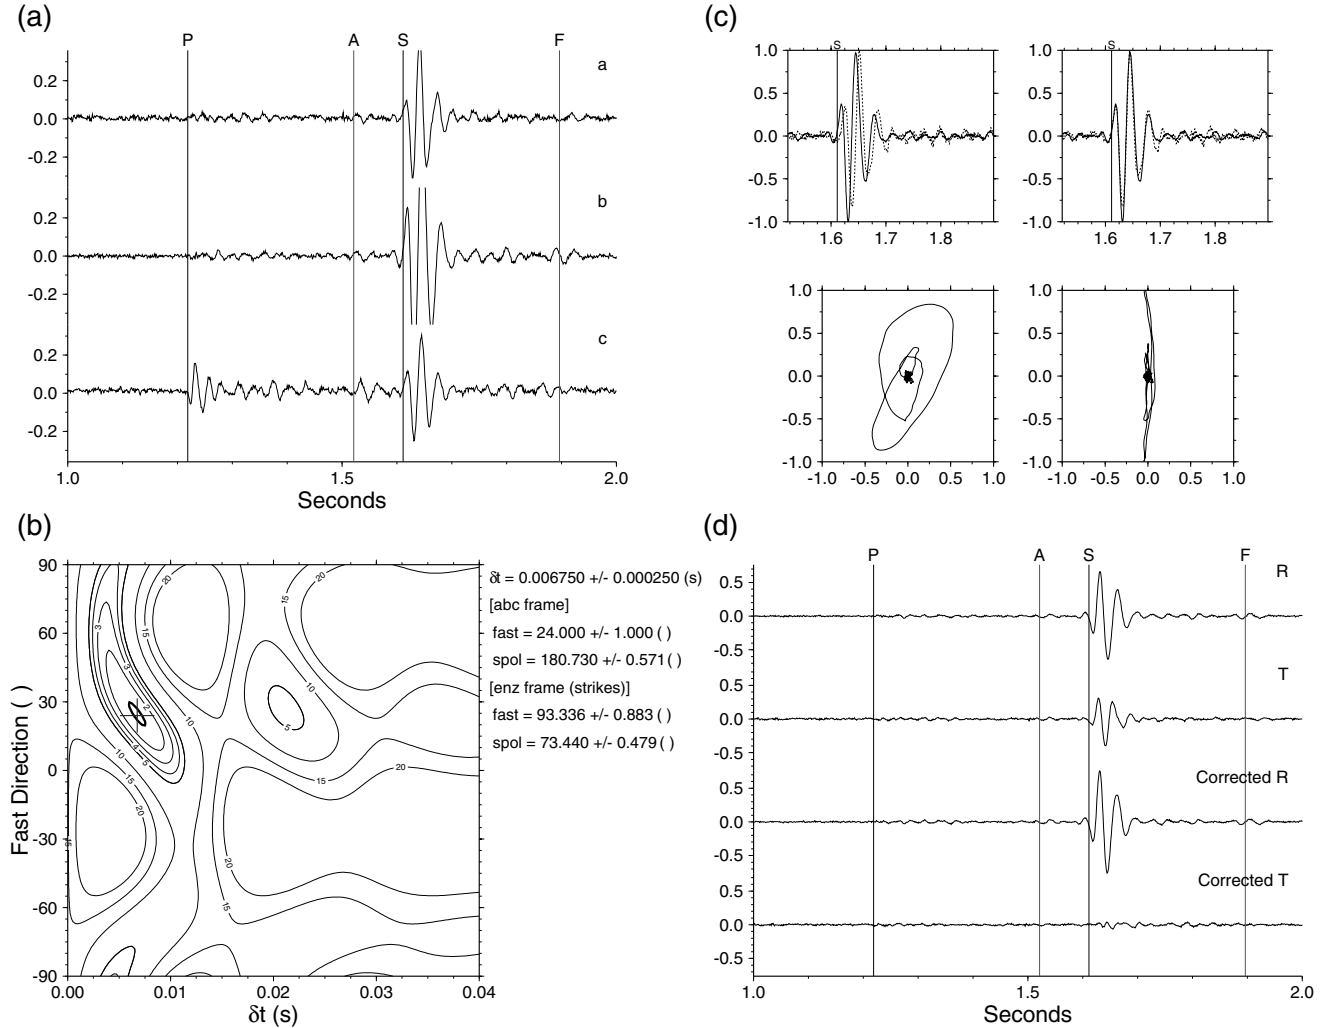


Figure 1. Illustration of Silver and Chan (1991) splitting correction technique on micro-seismic data. (a) Raw data rotated into the frame of the ray such that *a* and *b* are perpendicular to the ray path and *c* is along the ray path. *P*-wave particle motion was used to determine the ray direction in this case (Teanby *et al.*, 2003). *P*, *P*-wave pick; *S*, *S*-wave pick; *A* and *F* are the beginning and end of the shear-wave analysis window. (b) A grid search over ϕ and δt is performed to find the parameters that best linearize the particle motion. The cross shows the solution, and the thick contour is the 95% confidence interval. The labels next to the figure give δt , fast direction, and source polarization direction. Directions are given in the rotated (abc) frame and as projections into the geographic (enz) frame. (c) Fast and slow shear waveforms (top) and particle motion (bottom) before (left) and after (right) the shear-wave splitting correction. The fast and slow waves have similar waveforms, and the particle motion has been linearized after the correction. (d) Radial and transverse components before and after the splitting correction. The energy should be minimized on the corrected transverse component in the shear-wave analysis window. This event meets all the reliability criteria and is considered a reliable result.

has been decided upon, we use criteria based on the variance of the clusters and measurements to determine the optimum cluster and from this cluster the shear-wave analysis window corresponding to the measurement with the smallest errors on ϕ and δt .

Shear-Wave Analysis Windows

The first step in our procedure is to set up a grid of analysis windows. The beginning of the analysis window

T_{beg} is allowed to vary between T_{beg_0} and T_{beg_1} , with N_{beg} steps of ΔT_{beg} . Similarly, the end of the analysis window T_{end} is allowed to vary between T_{end_0} and T_{end_1} , with N_{end} steps of ΔT_{end} . The total number of analysis windows N is therefore given by $N = N_{\text{beg}}N_{\text{end}}$, and the shear-wave analysis window is defined by

$$T_{\text{beg}} = T_{\text{beg}_1} - (i - 1)\Delta T_{\text{beg}} \text{ for } i = 1 \dots N_{\text{beg}}, \quad (1)$$

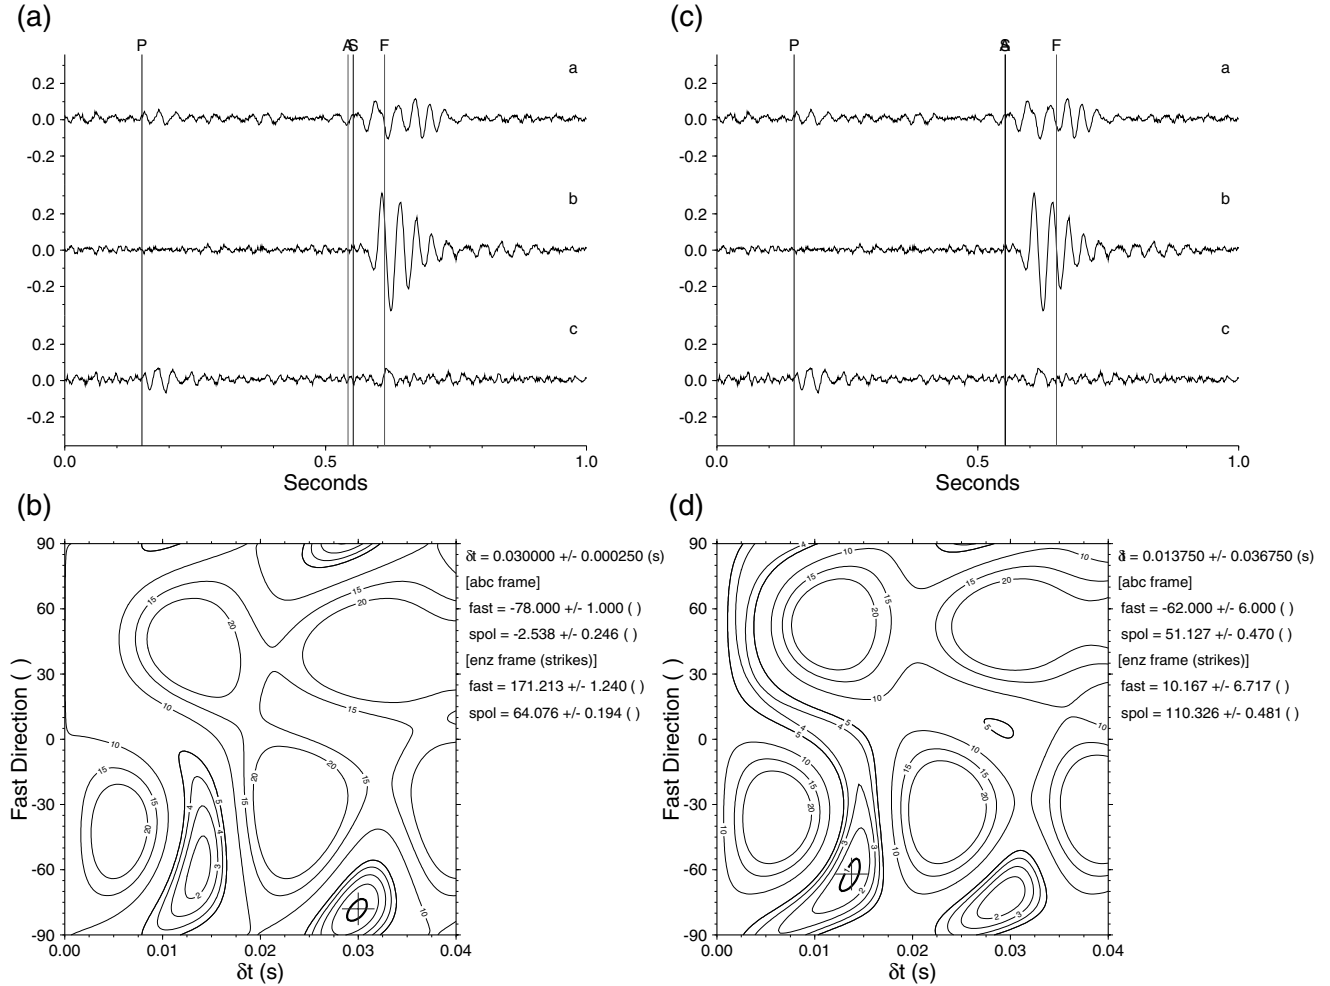


Figure 2. The effect of shear-wave analysis window choice on splitting measurements. (a) and (b) show the first choice of window and error surface. The solution is $\delta t = 30$ msec and $\phi = -78^\circ$. (c) and (d) show a second, slightly different choice of window and corresponding error surface. This time the solution is $\delta t = 13.75$ msec and $\phi = -62^\circ$, a difference of over 16 msec. Both solutions meet all the reliability criteria. This illustrates how slight changes in the analysis window can cause very different solutions. In this case, the difference is due to cycle skipping because of the band/limited nature of the data. However, it is not clear which one is the correct answer.

$$T_{\text{end}} = T_{\text{end}_0} + (j - 1)\Delta T_{\text{end}} \text{ for } j = 1 \dots N_{\text{end}}. \quad (2)$$

Each of the N analysis windows is assigned a unique window number equal to $(j - 1)N_{\text{beg}} + i$. However, this numbering is only used for creating plots like Figure 3; the ordering of windows is not important for the cluster analysis. Generally, larger window numbers correspond to longer windows. T_{beg_0} , T_{beg_1} , T_{end_0} , and T_{end_1} are all defined relative to the S -wave pick T_s . T_{beg_1} and T_{end_0} define the minimum analysis window. The extremes T_{beg_0} and T_{end_1} must be such that the window does not include any other phases.

Clustering and Stopping Criteria

Once the windows are defined, the splitting technique of Silver and Chan (1991) is applied for each window. This results in a set of N measurements of δt and ϕ . An example

of these measurements is shown in Figure 3 as a function of window number. We assume that reliable results are stable over many different windows. Therefore, robust measurements should be contained within clusters of similar measurements. Figure 4 shows a scatter plot of the 250 measurements of δt and ϕ shown in Figure 3; measurements fall into several distinct clusters. To identify these clusters, we use an unsupervised cluster analysis (see Everitt *et al.* [2001] for an overview of cluster analysis).

Implicit to all clustering techniques is the calculation of distances between data points and/or clusters. Because δt and ϕ are on different scales (0–40 msec and -90° to 90° , respectively) the data needs to be scaled or standardized before clustering to give equal weighting to each variable when assessing the proximity of clusters. Otherwise variations in δt will be judged less severely than variations in ϕ . We scale




Figure 3. Measurements of δt and ϕ obtained from 250 different analysis windows plotted against window number (window length tends to increase with window number). Stable solutions correspond to plateaus in this plot with small error bars. The solution that was selected using cluster analysis is marked with the gray crosshairs. Note the large jumps in both ϕ and δt as window number increases.




Figure 4. Measurements of δt and ϕ obtained from 250 different analysis windows. The plateaus from Figure 3 condense into tight clusters of points. Many points in the clusters lie on top of each other because δt and ϕ are found using a grid search. Clusters are automatically identified using cluster analysis, with triangles marking the cluster positions and a cross marking the optimum cluster, with the lowest variance σ_o^2 .

using the range of the variables, in our case 180° and 40 msec. Scaling by variable range has been shown to perform very well in many clustering applications (Milligan and Cooper, 1988; Everitt *et al.*, 2001). From now on, δt and ϕ will refer to the rescaled measurements.

Consider the N scaled measurements $(\delta t_i, \phi_i)$ with variances $(\sigma_{\delta t_i}^2, \sigma_{\phi_i}^2)$, $i = 1 \dots N$. The data will be partitioned into M clusters. In each cluster C_j , there are N_j data points, where $j = 1 \dots M$. The δt_i and ϕ_i data points are clustered using an hierarchical technique (Everitt *et al.*, 2001). We start with the same number of clusters as there are data points ($M = N$) and calculate all the intercluster distances, taking into account the cyclic nature of ϕ . The intercluster distance is simply the rescaled Euclidian distance between the cluster centers. In using this distance measure, we assume that clusters are characterized by isotropic variances, which is appropriate for our application as we have no *a priori* knowledge of the cluster structure. The two nearest clusters are then combined so that the number of clusters decreases by one. We continue combining clusters until there is only one cluster ($M = 1$) comprising the whole dataset. The result is a hierarchy of clusters.

For each number of clusters $M = 1 \dots N$, we calculate the number of data points N_j in each cluster C_j and the positions of the cluster centers $(\Delta t_j, \Phi_j)$, given by the mean position of points within the cluster:

$$\Delta t_j = \frac{\sum_{i=1}^{N_j} \delta t_i^{(j)}}{N_j}, \quad (3)$$

$$\Phi_j = \frac{\sum_{i=1}^{N_j} \phi_i^{(j)}}{N_j}, \quad (4)$$

where $\phi_i^{(j)}$ and $\delta t_i^{(j)}$ are the measurements that belong to cluster j . The final clustering requires that we define the number of clusters M . We wish to do this in an unsupervised manner, that is, without any need for user input. This is a difficult problem in cluster analysis, and there are many methods for assessing the optimum number (Milligan and Cooper, 1985). Fortunately the number of clusters is not critical for our analysis, as stable plateau regions correspond to fairly tight clusters. The hierarchical clustering methodology will identify very tight clusters first, and these should have a large number of points. To determine the number of clusters, we use the methods of Caliński and Harabasz (1974) and Duda and Hart (1973). These were the top two performers in a comparison of 30 estimators of optimum cluster number by Milligan and Cooper (1985). Clustering is stopped when these criteria pass specific thresholds.

We define the within-cluster covariance \mathbf{W} and the between-cluster covariance \mathbf{B} as

$$\mathbf{W} = \begin{pmatrix} \sum_{j=1}^M \sum_{i=1}^{N_j} (\delta t_i^{(j)} - \Delta t_j)^2 & \sum_{j=1}^M \sum_{i=1}^{N_j} (\delta t_i^{(j)} - \Delta t_j)(\phi_i^{(j)} - \Phi_j) \\ \sum_{j=1}^M \sum_{i=1}^{N_j} (\delta t_i^{(j)} - \Delta t_j)(\phi_i^{(j)} - \Phi_j) & \sum_{j=1}^M \sum_{i=1}^{N_j} (\phi_i^{(j)} - \Phi_j)^2 \end{pmatrix}, \quad (5)$$

$$\mathbf{B} = \begin{pmatrix} \sum_{j=1}^M (\Delta t_j - \bar{\Delta t})^2 & \sum_{j=1}^M (\Delta t_j - \bar{\Delta t})(\Phi_j - \bar{\Phi}) \\ \sum_{j=1}^M (\Delta t_j - \bar{\Delta t})(\Phi_j - \bar{\Phi}) & \sum_{j=1}^M (\Phi_j - \bar{\Phi})^2 \end{pmatrix}, \quad (6)$$

where $\bar{\Delta t}$ and $\bar{\Phi}$ are the mean values of δt and ϕ over all the samples:

$$\bar{\Delta t} = \frac{\sum_{i=1}^N \delta t_i}{N}, \quad (7)$$

$$\bar{\Phi} = \frac{\sum_{i=1}^N \phi_i}{N}. \quad (8)$$

The Caliński and Harabasz (1974) criterion is

$$c(M) = \frac{(N - M)\text{trace}(\mathbf{B})}{(M - 1)\text{trace}(\mathbf{W})}. \quad (9)$$

The maximum of this function indicates that the between-cluster variance is maximized with respect to the within-cluster variance, indicating tight clusters that are widely spaced. Therefore, the optimum number of clusters M is obtained when $c(M)$ is maximized.

The Duda and Hart (1973) criterion is based on the ratio of within-cluster variances when two clusters are combined into one cluster. The variance of the two individual clusters is given by

$$\sigma_2^2 = \sum_{j=1}^2 \sum_{i=1}^{N_j} [(\delta t_i^{(j)} - \Delta t_j)^2 + (\phi_i^{(j)} - \Phi_j)^2], \quad (10)$$

and the variance when the two clusters are combined into one cluster is

$$\sigma_1^2 = \sum_{i=1}^{N_1} [(\delta t_i^{(1)} - \Delta t_1)^2 + (\phi_i^{(1)} - \Phi_1)^2]. \quad (11)$$

The null hypothesis is that the two clusters should be combined as a single cluster. Normally distributed within-cluster distances are assumed, and the null hypothesis is rejected when

$$\left(1 - \frac{\sigma_2^2}{\sigma_1^2} - \frac{2}{\pi p}\right) \left(\frac{N_j p}{2[1 - 8/(\pi^2 p)]}\right)^{1/2} > c_{\text{critical}}, \quad (12)$$

where p is the number of parameters ($= 2$, i.e., ϕ and δt) and c_{critical} is the critical value from a standard normal distribu-

tion. Milligan and Cooper (1985) found that $c_{\text{critical}} = 3.20$ gave the best results. We consider the hierarchy of clusters from $M = 1 \dots N$ and halt the subdivision of clusters when equation (12) is no longer satisfied.

We used the maximum value of M predicted by the two stopping criteria as the optimum number of clusters because in our case it is preferable to overestimate the number of clusters so that significantly different results are not included in the same cluster. The preceding criteria given for determining M can become unstable for high numbers of clusters. Therefore, it is advisable to set an upper limit on the number of clusters M_{max} . As there should be relatively few clusters, the choice of this upper limit is not critical. Three clusters were found in Figure 4.

Selection of Optimum Cluster and Measurement

Once the cluster centers and optimum number of clusters have been determined, we must select the best cluster and the best measurement from within this cluster. Criteria for the best cluster are based on the number of points and the variance within the cluster. All clusters with less than $N_{c_{\text{min}}}$ data points are considered spurious and rejected. If this leaves no clusters, then there is no stable solution. $N_{c_{\text{min}}}$ is

chosen such that it corresponds to approximately a cycle's worth of points.

The within-cluster variance $\sigma_{c_j}^2$ and mean data variance, $\sigma_{d_j}^2$ of the remaining clusters are then calculated according to

$$\sigma_{c_j}^2 = \frac{\sum_{i=1}^{N_j} (\delta t_i^{(j)} - \Delta t_j)^2 + (\phi_i^{(j)} - \Phi_j)^2}{N_j}, \quad (13)$$

$$\sigma_{d_j}^2 = \left[\sum_{i=1}^{N_j} \frac{1}{(\sigma_{\delta t_i}^2)^2} \right]^{-1} + \left[\sum_{i=1}^{N_j} \frac{1}{(\sigma_{\phi_i}^2)^2} \right]^{-1}. \quad (14)$$

Equation (14) is related to the harmonic mean, which reduces the effect of outliers. Figure 3 shows that there are stable regions (tight clusters) with large measurement errors, especially for long windows (high window numbers), characterized by a high value of $\sigma_{d_j}^2$. It is also possible to have unstable measurements (diffuse clusters) with low measurement errors, especially for short windows, characterized by a high value of $\sigma_{c_j}^2$. Therefore, we define an overall variance for the cluster $\sigma_{o_j}^2$, which is set to $\max(\sigma_{c_j}^2, \sigma_{d_j}^2)$. The best cluster has the smallest value of $\sigma_{o_j}^2$. Although $\sigma_{o_j}^2$ is a simple measure of cluster quality, minimizing $\sigma_{o_j}^2$ avoids the selection of diffuse clusters with low measurement errors and tight clusters with high measurement errors, as illustrated in Figure 5.

The best measurement is simply the measurement with the smallest variance from within the best cluster. The best measurement from the best cluster in our example is shown with crosses in Figures 3 and 4. The splitting results using this window are shown in Figure 6.

Diagnostic Plot

The diagnostic plot shown in Figure 6 is produced in order to evaluate if the automated method has produced a reliable measurement. The result is considered reliable if (1) the energy on the corrected transverse component is minimized and about the same level as the noise, (2) the fast and slow waveforms are similar and match well after correction, (3) the particle motion in the shear-wave analysis window is elliptical before correction and linear after, and (4) the error surface has a unique well-constrained solution (i.e., no cycle skipping and a small 95% confidence contour). The scatter plot of δt and ϕ should be inspected to ensure that the number and positions of the automatically determined clusters are reasonable. The stability of the splitting result can be assessed by inspecting the variation of δt and ϕ with window number (e.g., Fig. 3). A stable splitting result should have one plateau containing many points with low errors.

Application to Microseismic Data

The method has been tested on a dataset of 324 located and picked microseismic events from the Valhall oil field in the central North Sea (Dyer and Jones, 1998). This is a chalk




Figure 5. Three different cluster types: (a) diffuse cluster with low individual measurement errors, (b) tight cluster with high individual measurement errors, and (c) tight cluster with low individual measurement errors. Cluster (a) has high σ_c^2 and low σ_d^2 , cluster (b) has low σ_c^2 and high σ_d^2 , and cluster (c) has low σ_c^2 and σ_d^2 . Cluster (c) has the lowest value of σ_o^2 and is the best cluster overall. We propose σ_o^2 as a measure of cluster quality.

reservoir overlain by siltstone caprock in which oil production has caused subsidence, leading to the occurrence of microseismic events as a mechanism for releasing stress. The microseismic events were recorded over an 8-week period on six three-component receivers in a vertical array located between 2000 and 2100 m depth in a disused borehole. Data from this microseismic study are characterized by lower noise levels than teleseismic studies because the geophones were buried at a depth of 2 km. We have previously performed shear-wave splitting analysis on this dataset using the method of Silver and Chan (1991) by manually selecting the shear-wave analysis windows (Teaby *et al.*, 2003). Prior to the splitting analysis, seismograms were rotated into the ray frame, using the *P*-wave particle motion as a proxy for ray direction, in order to maximize the *S*-wave energy on the two components used in the splitting analysis. Over 100 reliable results were obtained for each of the six receivers, exhibiting average splitting of around 10 msec. The large number of events and availability of previous reliable splitting measurements makes this dataset an excellent test case for our automated technique.

We applied the automated method to the subset of microseismic data from Teaby *et al.* (2003) that gave reliable shear-wave splitting measurements using manual windowing, so that manual and automated results could be compared. The parameters used are given in Table 1. Parameters were chosen such that the shear-wave analysis windows (1) did not overlap with the *P* wave, which preceded the *S* wave by well over 0.1 sec; (2) could encompass all of the shear-wave energy envelope, which was typically less than 0.25 sec; and (3) included a minimum of one wavelength of the dominant period (≈ 0.04 sec at 30 Hz). $N_{c_{min}}$ was chosen such that clusters contained at least a number of points equivalent to a variation in window length of one dominant period.

Figure 7 compares the manually windowed results from Teaby *et al.* (2003) and the results obtained from the au-

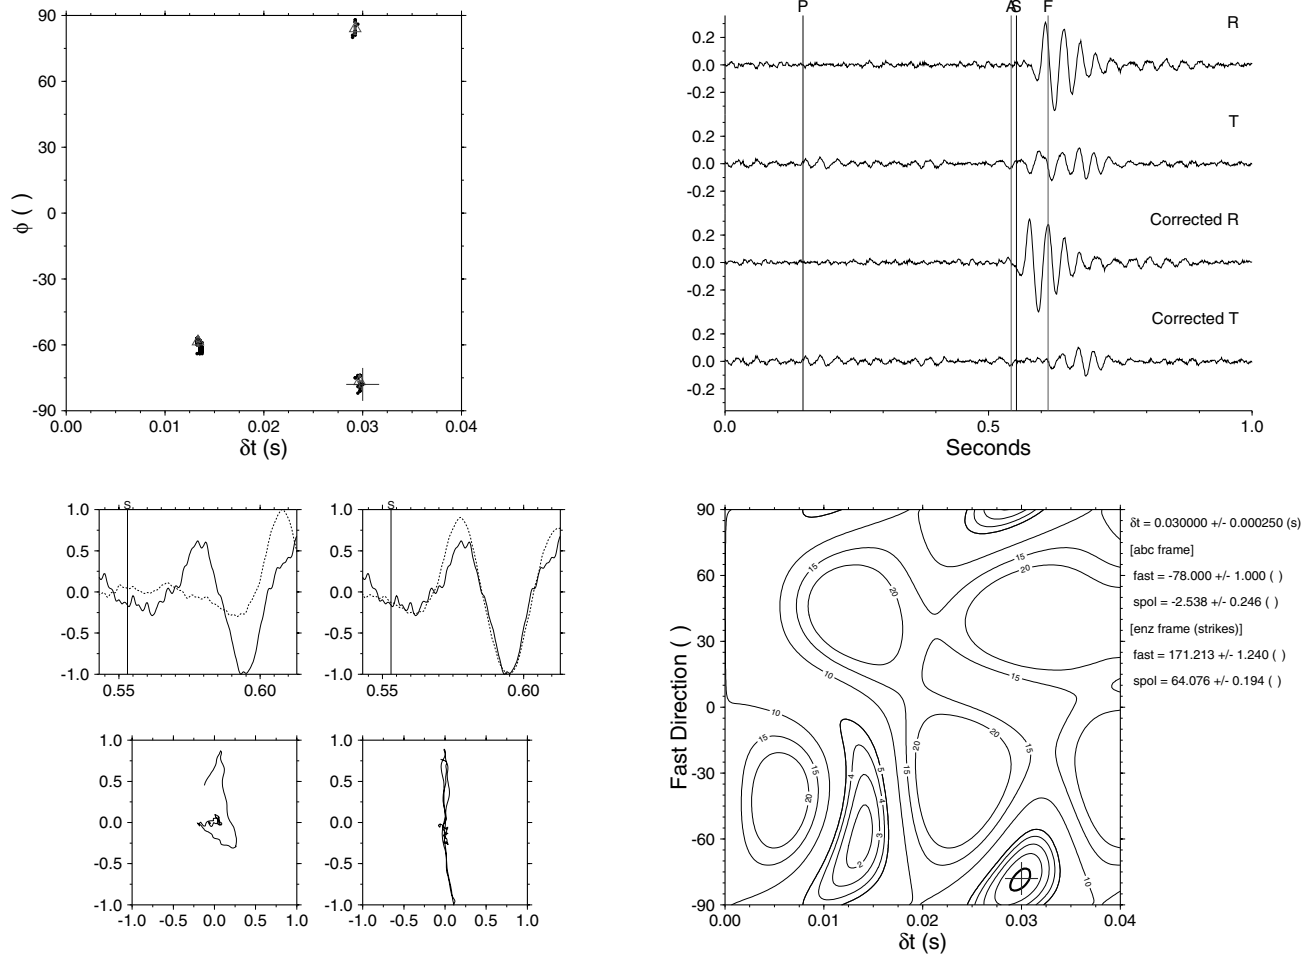


Figure 6. Diagnostic plot used to determine if the automated result is reliable. For a reliable result, (1) the energy on the corrected transverse component should be minimized, (2) the fast and slow shear waves should be similar in shape and the particle motion should be linearized after the splitting correction, (3) the error surface should have a well constrained and unique solution, and (4) the cluster analysis has identified the three clusters correctly. This example passes all the criteria and is considered reliable.

Table 1

Parameters used for Automated Splitting Measurements on the Valhall Microseismic Dataset

Parameter	Value
N_{beg}	5
N_{end}	50
N	$250 (= N_{\text{beg}} \times N_{\text{end}})$
ΔT_{beg}	0.020 sec
ΔT_{end}	0.005 sec
T_{beg_1}	0.010 sec before T_s
T_{end_0}	0.040 sec after T_s
$N_{\text{c}_{\text{min}}}$	10
M_{max}	20

tomated method. The results shown are for receiver number 1 (top of the array); the other five receivers give similar results. Figure 7a,b shows all the results. There is a one-one relationship between most of the manual and automated results. However, there is quite a lot of scatter about this line. The scatter is much reduced in Figure 7c,d, where results with large error bars (over 15° in ϕ or over 1.5 msec in δt) have been rejected. In Figure 7e,f, the diagnostic plots (like the plot shown in Fig. 6) have been examined, and solutions that do not satisfy the requirements for reliable splitting measurements specified in the Diagnostic Plot section have been rejected. This manual quality control is an important step and has significantly improved the agreement between the manual and automated measurements.




Figure 7. Comparison of automated and manually windowed splitting results on the Valhall microseismic dataset. The percentage of measurements where manual and automatic results agree to within 15° or 1.5 msec is given in square brackets. (a) and (b) show all the results, with no user interference or grade selection [64%]. (c) and (d) show results after screening such that error bars are less than 15° or 1.5 msec [82%]. In (e) and (f) the results have been manually screened by inspecting the diagnostic plots [92%]. There is very good agreement between manual and automatic results, especially after rejecting measurements with large errors and manual quality control.

Discussion

In order to choose the window parameters, we must consider (1) the dominant period of the S wave, (2) separation of the S phase from other phases, and (3) the duration of S -wave energy envelope. The most computationally important parameters are N_{beg} , N_{end} , Δt_{beg} , and Δt_{end} . Large N and small Δt give the most detailed exploration of the window space. However, computation time is directly proportional to N , so a compromise is required. Splitting measurements are usually more sensitive to the window end than the window beginning, so N_{beg} can be smaller than N_{end} and Δt_{beg} can be larger than Δt_{end} . The choice of Δt_{end} depends on the sensitivity of the splitting measurement to window choice, but we found that a Δt_{end} equivalent to about a tenth of a cycle gave more than enough detail in most cases. T_{beg_1} and T_{end_0} , which define the minimum window, are chosen such that the minimum window starts slightly before the S -wave pick, by an amount greater than the S -wave pick uncertainty, and the window ends at least half a cycle after T_s . This cor-

responds to the minimum acceptable windows used in manual analyses.

Parameter values should be chosen such that the extremes of the analysis window, T_{beg_0} and T_{end_1} , do not overlap with other phases that have different amounts of splitting. Although the choice of T_{beg_0} and T_{end_1} is subjective, the results will not be affected if all of the S -wave energy envelope is included. However, when neighboring phases occur very close together, it becomes difficult to avoid including interfering phases in the analysis window. Here the choice of T_{beg_0} and T_{end_1} is more subjective, but the final splitting measurement should still be robust as long as a large enough range of analysis windows are included. If multiple phases, with different amounts of splitting, are included in the analysis window, the automated method will either find no stable solution, a spurious unreliable solution, or multiple clusters with splitting parameters corresponding to each phase. However, all these cases should show up on the diagnostic plot and some useful information may still be extracted, although

this would complicate the analysis and is best avoided if possible.

The maximum number of clusters M_{\max} is not critical, as there are usually few clusters for well-constrained results. A value of 20 should be applicable to most shear-wave splitting studies and stop spuriously high numbers of clusters being defined. The minimum number of measurements N_{\min} required for an acceptable cluster should be equivalent to about a wavelength's worth of points. However, it is advisable for N_{\min} to be less than N/M_{\max} , especially if an initial coarse grid search is being performed.

The method leads to code that is easily parallelized, which makes the code simple to implement on parallel processors. It also makes the choice of compiler and optimization important. For the Valhall dataset, each event took 15 sec to process on a single 2-GHz processor with $N = 250$ and using a fully optimized compiler. For an experienced interpreter, the inspection of diagnostic plots is very fast (e.g., quality control on 152 events took 1 hr). This compares well with manual window selection, which took around 1 week of solid work to study the same 152 events.

Of the 152 events processed for receiver 1, which all had reliable splitting measurements when the analysis window was manually selected, 91 of the automated results passed the quality control after inspection of the diagnostic plots. Therefore, using the automated technique resulted in about 40% fewer measurements than for manual window picking. Some of these losses are due to rejection of events that had window-sensitive results. However, an experienced interpreter can always extract measurements from more events than an automated technique. The loss of some suitable events is a price we need to pay for automated analyses. On the other hand, an automated approach can handle data volumes that would otherwise be prohibitively labor intensive. Hence, the loss of events is not a significant problem for large datasets.

To test if the method was generally applicable, we applied it to teleseismic S phases recorded in Australia, previously published by Wookey *et al.* (2002). The grid search covered δt 's of up to 4 sec, and the parameters used in the automation were the same as in Table 1 except $\Delta T_{\text{beg}} = 2.5$ sec, $\Delta T_{\text{end}} = 0.5$ sec, $T_{\text{beg}_1} = 1$ sec, and $T_{\text{end}_0} = 8$ sec. Initial application to filtered teleseismic S -wave data was successful and agreed with the manually windowed results from Wookey *et al.* (2002). Even though each event was studied individually, the method still speeded up analysis by automating the choice of window within a user-defined maximum window, thus allowing us to assess the stability of the measurement without resorting to labor-intensive repeated manual analyses. Results could also be considered more robust as uncertainties about window dependence were removed. Error bars on splitting measurements were reduced due to optimization of the analysis window.

For this small teleseismic S -wave dataset (and also the microseismic data), we used the same windowing parameters for the whole dataset. In a general teleseismic dataset,

there is a large range of source-receiver distances, so some parameters must be varied with distance. The maximum δt in the error surface grid search could be set to be some fraction of the total travel time from source to receiver. This will give a more appropriate scaling of the δt measurements for the cluster analysis. Travel-time tables could be used to predict the arrival of different phases and define the windowing parameters so as to avoid including multiple phases in the same analysis window. Additionally, a wave-field decomposition method similar to that used by Kennett (1991) and Bostock (1998) could be used to separate P and S phases.

Examination of the results from the Valhall microseismic and Australian teleseismic data showed that the number of clusters can vary from a single cluster, for very well constrained measurements, to many clusters for less well constrained measurements. Cluster shapes were generally circular for well-constrained measurements, so that our assumption of isotropic variance was valid. However, for small amounts of splitting δt , the fast direction ϕ is poorly constrained and clusters tend to be elongated in the ϕ direction. For these cases, the cluster analysis tended to split the cluster into several small circular clusters. This could be overcome by rescaling the variables over a different range, thereby using anisotropic variances. Alternatively, the distance between closest points in two clusters could be used as the distance measure instead of the cluster centers (Everitt *et al.*, 2001). However, we did not use this method, as it tends to underestimate the number of clusters.

As mentioned in the Method section and shown in Figure 2, the strong window dependence/cycle skipping of some splitting results is a serious problem. This problem is most severe for band-limited data. Using an automated technique reduces subjectivity of measurements but does not solve this cycle-skipping problem. Wolfe and Silver (1998) discussed a method to stack the error surfaces from multiple events/stations for poorly constrained splitting measurements. This adds consistency of measurements as an extra constraint and helps to remove spurious results. The technique of Wolfe and Silver (1998) could be applied directly to combine the individual automated results, or alternatively a cluster analysis could be done on multiple events/stations simultaneously. Also, the results from the first and second best cluster could be compared. If the first is obviously better than the second then the result is reliable, otherwise results may be affected by cycle skipping.

An important limitation of this cluster-based automation method is that it cannot distinguish between a null measurement, where there is no splitting, and a poorly constrained result from noisy data. This is because the method relies on identifying tight clusters of measurements using cluster analysis. It is therefore applicable to seismograms displaying finite anisotropy, where tight clusters form. However, for null measurements ϕ is totally unconstrained, leading to a large spread in δt and an elongated 95% confidence contour. Hence null measurements tend to have a large scatter of points on a plot of δt versus ϕ ; consequently, no best cluster

can be identified and the measurement is often rejected. However, unlike noisy results, null measurements have linear particle motion and well-matched fast and slow waves before any splitting correction is applied, which could be identified by manual inspection of diagnostic plots. So a semi-automatic null detection should be possible but would require an extra quality-control step.

Conclusions

We have developed a method to automate shear-wave splitting analysis. The most important features of this method are that it removes the subjectivity of window selection and increases the speed of data analysis. The method can be used to improve the quality of shear-wave splitting measurements and can be applied to large datasets that would otherwise require a prohibitively long manual analysis. The method requires travel-time picks for the *S* phase to be studied and a set of windowing parameters such that interfering phases are not included in the analysis windows. A limitation of the method is that null results are not automatically detected. However, an extra quality-control step may enable better detection.

Manual quality control of results is recommended. However, this step is significantly faster than a complete manual analysis, especially when the choice of windows is not obvious. A diagnostic plot shows quickly whether the results are satisfactory. The method has been applied to both microseismic and teleseismic events and works well in both cases.

Acknowledgments

This work was supported by ABB Offshore Systems, BP Norge, Shell Exploration and Production UK, Schlumberger Cambridge Research, and NERC (under the Leeds Anisotropy and Microseismicity Project). The authors thank O. Barkved (BP) and R. H. Jones (ABB) for providing the Valhall dataset and G. Helffrich and M. Savage for providing insightful reviews that helped to improve the manuscript.

References

- Aster, R. C., P. M. Shearer, and J. Berger (1990). Quantitative measurements of shear-wave polarizations at the Anza seismic network, southern California: implications for shear-wave splitting and earthquake prediction, *J. Geophys. Res.* **95**, 12,449–12,473.
- Bokelmann, G. H. R., and H. P. Harjes (2000). Evidence for temporal variation of seismic velocity within the upper continental crust, *J. Geophys. Res.* **105**, 23,879–23,894.
- Bostock, M. G. (1998). Mantle stratigraphy and evolution of the Slave province, *J. Geophys. Res.* **103**, 21,183–21,200.
- Calinski, T., and J. Harabasz (1974). A dendrite method for cluster analysis, *Comm. Stati.* **3**, 1–27.
- Duda, R. O., and P. E. Hart (1973). *Pattern Classification and Scene Analysis*, Wiley, New York.
- Dyer, B., and R. Jones (1998). 3D processing and interpretation of the microseismic data recorded during June and July 1998 in the Valhall field, unpublished report for Amoco by CSMA Ltd. Cornwall.
- Everitt, B. S., S. Landau, and M. Leese (2001). *Cluster Analysis*, Fourth Ed., Arnold, London.
- Gledhill, K., and G. Stuart (1996). Seismic anisotropy in the fore-arc region of the Hikurangi subduction zone, New Zealand, *Phys. Earth Planet. Interiors* **95**, 211–225.
- Kendall, J.-M. (2000). Seismic anisotropy in the boundary layers of the mantle, in *Earth's Deep Interior: Mineral Physics and Tomography from the Atomic to the Global Scale*, S. Karato, A. R. L. Forte, G. Masters, and L. Stixrude (Editors), American Geophysical Monograph 117, 133–159.
- Kennett, B. L. N. (1991). The removal of free surface interactions from three-component seismograms, *Geophys. J. Int.* **104**, 153–163.
- Milligan, G. W., and M. C. Cooper (1985). An examination of procedures for determining the number of clusters in a data set, *Psychometrika* **50**, 159–179.
- Milligan, G. W., and M. C. Cooper (1988). A study of standardization of variables in cluster-analysis, *J. Classification* **5**, 181–204.
- Park, J., and V. Levin (2002). Seismic anisotropy: tracing plate dynamics in the mantle, *Science* **296**, 485–489.
- Sandvol, E., and T. Hearn (1994). Bootstrapping shear-wave splitting errors, *Bull. Seism. Soc. Am.* **84**, 1971–1977.
- Savage, M. K. (1999). Seismic anisotropy and mantle deformation: what have we learned from shear wave splitting? *Rev. Geophys.* **37**, 65–106.
- Savage, M. K., X. R. Shih, R. P. Meyer, and R. C. Aster (1989). Shear-wave anisotropy of active tectonic regions via automated S-wave polarization analysis, *Tectonophysics* **165**, 279–292.
- Silver, P. G. (1996). Seismic anisotropy beneath the continents: probing the depths of geology, *Ann. Rev. Earth Planet. Sci.* **24**, 385–432.
- Silver, P. G., and W. W. J. Chan (1991). Shear-wave splitting and subcontinental mantle deformation, *J. Geophys. Res.* **96**, 16,429–16,454.
- Teanby, N., J.-M. Kendall, R. H. Jones, and O. Barkved (2004). Stress-induced temporal variations in seismic anisotropy observed in microseismic data, *Geophys. J. Int.* (in press).
- Vinnik, L. P., R. Kind, G. L. Kosarev, and L. I. Makeyeva (1989). Azimuthal anisotropy in the lithosphere from observations of long-period S-waves, *Geophys. J. Int.* **99**, 549–559.
- Wolfe, C. J., and P. G. Silver (1998). Seismic anisotropy of oceanic upper mantle: shear-wave splitting methodologies and observations, *J. Geophys. Res.* **103**, 749–771.
- Wookey, J., J. M. Kendall, and G. Barruol (2002). Mid-mantle deformation inferred from seismic anisotropy, *Nature* **415**, 777–780.

School of Earth Sciences
University of Leeds
Leeds, LS2 9JT
United Kingdom

Manuscript received 17 June 2003.