JOURNAL OF GEOPHYSICAL RESEARCH, VOL. 112, A04204, doi:10.1029/2006JA011900, 2007

An automatic wave detection algorithm applied to Pc1 pulsations

J. Bortnik,[1,3] J. W. Cutler,[2,3] C. Dunson,[4] and T. E. Bleier[4]

Received 7 June 2006; revised 7 November 2006; accepted 29 November 2006; published 6 April 2007.

[1] A new technique designed to automatically identify and characterize waves in three-axis data is presented, which can be applied in a variety of settings, including triaxial ground-magnetometer data or satellite wave data (particularly when transformed to a field-aligned coordinate system). This technique is demonstrated on a single Pc1 event recorded on a triaxial search coil magnetometer in Parkfield, California (35.945°,−120.542°), and then applied to a 6-month period between 1 June 2003 and 31 December 2003. The technique begins with the creation of a standard dynamic spectrogram and consists of three steps: (1) for every column of the spectrogram (which represents the spectral content of a short period in the time series), spectral peaks are identified whose power content significantly exceeds the ambient noise; (2) the series of spectral peaks from step 1 are grouped into continuous blocks representing discrete wave events using a ''spectral-overlap'' criterion; and (3) for each identified event, wave parameters (e.g., wave normal angles, polarization ratio) are calculated which can be used to check the continuity of individual identified wave events or to further filter wave events (e.g., by polarization ratio).

Citation: Bortnik, J., J. W. Cutler, C. Dunson, and T. E. Bleier (2007), An automatic wave detection algorithm applied to Pc1 pulsations, *J. Geophys. Res.*, *112*, A04204, doi:10.1029/2006JA011900.

1. Introduction

[2] A large variety of plasma waves are routinely encountered in studies of the magnetosphere, carrying with them information about the generation mechanism and medium of propagation between source and receiver [e.g., *Gurnett and Inan*, 1988; *Samson*, 1991; *Stix*, 1992]. Such waves have been studied for well over 100 years [*Mursula et al.*, 1994] and have provided invaluable information about the space environment.

[3] Currently, there exist large quantities of wave data gathered from over 50 years of space and ground observations [*Walker et al.*, 2005], which are growing daily due to the unprecedented number of deployed instruments and affordability of mass storage. The exponential increase in computational speed [e.g., *Moore*, 1965] allows analysis of these large data sets in reasonable periods of time and only requires suitable algorithms to extract information about the waves. Such information would ideally include the shape of the wave event in the frequency-time domain (specified as upper- and lower-frequency bounds which change as a function of time throughout the event), intensity, orientation of the wave normal, as well as wave polarization parameters such as polarization ratio, ellipticity, and sense of rotation.

[4] Previously, wave events in the Pc1 frequency range (0.2–5 Hz [*Jacobs*, 1970, p. 20]) were identified either manually [e.g., *Campbell and Stiltner*, 1965; *Fraser*, 1968; *Anderson et al.*, 1992a, 1992b; *Fraser and Nguyen*, 2001; *Meredith et al.*, 2003] by examining a series of spectrograms or with a variety of simple automated routines. For example, *Anderson et al.* [1992a] used spectrograms of their data and examined each column for a (single) five-point sliding average peak which sufficiently exceeded a threshold power level. Similarly, *Erlandson and Anderson* [1996] treated each column of the spectrogram (each column representing ~30 s) as an individual wave event and searched for (multiple) spectral peaks exceeding a threshold level. *Loto'aniu et al.* [2005] used a threshold in both duration and intensity in electric and magnetic components of the wave individually to identify wave events. While the above algorithms identify the presence of wave events, they do not extract the frequency-time ($f - t$) shape of the event. Such information could be obtained using an edge-detection [*Canny*, 1986] algorithm applied to the spectrogram but this approach does not (by itself) yield any further wave parameters. Other approaches (VLF range) have used matched filtering to extract wave information [*Hamar and Tarcsai*, 1982; *Singh et al.*, 1999] but require a priori knowledge of the shape of the spectrum, which limits their generality.

[5] In the present paper we introduce a simple technique which is nevertheless useful in automatically detecting and characterizing wave events in time series data. This tech-

[1]Department of Atmospheric and Oceanic Sciences, University of California, Los Angeles, California, USA.
[2]Department of Aeronautics and Astronautics, Space Systems Development Laboratory, Stanford University, Palo Alto, California, USA.
[3]Also at QuakeFinder, LLC., Palo Alto, California, USA.
[4]QuakeFinder, LLC., Palo Alto, California, USA.

0148-0227/07/2006JA011900$09.00

nique is presented in the context of detecting a typical Pc1 pulsation in three-component magnetometer data but can be applied generally to a variety of wave types and frequency regimes, for waves recorded using ground instruments or on satellites (particularly when the satellite data is rotated into a field-aligned coordinate system). In section 2 we present our methodology, which consists of identifying spectral peaks in short segments of time series data, temporally grouping the spectral peaks into blocks, and finally extracting useful wave characteristics for each block (e.g., polarization ratio, ellipticity, and wave normal direction). In section 3 we apply our algorithm to 6 months of magnetometer data and discuss various aspects and extensions of our technique using this example data. In section 4 we summarize our technique and findings.

2. Methodology

[6] We use the Fast Fourier Transform (FFT) implemented in a conventional dynamic spectrogram as the starting point for our identification algorithm [*Bracewell*, 2000, p. 491]. This is done for a number of reasons: (1) dynamic spectrograms are a common method of analyzing wave phenomena and are routinely generated. Thus it is expedient for us to use the output of the dynamic spectrograms as a starting point for our detection algorithm; (2) after the particular time series has been converted into a dynamic spectrogram format, our technique proceeds in a standard way through the analysis, regardless of the frequency range or type of wave under consideration, making this technique fairly general and independent of the type of wave and platform upon which it is measured; (3) the information obtained from dynamic spectrograms is particularly useful in that we can directly calculate a number of wave parameters from it such as polarization ratio, ellipticity, and wave normal orientation. These wave parameters can be used as a further check on the identification algorithm.

[7] In the subsections that follow, we use as an example the magnetic field data from a triaxial search coil magnetometer at Parkfield, California (Geographic: (35.945°, −120.542°), CGM: (41.61°, −56.8°), dip: 60.2°, declination: 14.7°, L value: 1.77), on 6 June 2003. This example contains a typical Pc1 wave event and is used to demonstrate our procedure which consists of three broad steps: (1) for every column of the spectrogram (which represents the spectral content of a short time period), spectral peaks are identified whose power content is significantly higher than the ambient noise (section 2.1); (2) given a series of spectral peaks from step 1, we temporally group the peaks into continuous blocks representing discrete wave events (section 2.2); and (3) for each identified wave event, polarization parameters are calculated which serve as a filter for either continuity or wave quality (section 2.3). At every step, there are a number of free parameters which can be chosen for the specific application at hand, and these parameters are discussed in the context of our example in the section below.

2.1. Frequency Band Identification

[8] We begin by creating a dynamic spectrogram of a long time series of sampled data in the usual way [*Bracewell*, 2000, p. 491]. In our case three time series of magnetic field intensity are used, representing each component of our triaxial magnetometer set. The data are sampled at f_s = 40 Hz and are processed in blocks of 1 day (3,456,000 samples). The time series are then divided into consecutive and overlapping time segments, each time segment is multiplied by a Hamming window to reduce edge effects, and the FFT is applied to the resulting time series. In our case, we have chosen the time segment to be $N_{ch} = 4096$ samples long (∼102.4 s per time segment), with an overlap of w_{ol} ∼30% (∼30.7 s), resulting in 1205 time segments per day with a ∼71.7 s spacing between adjacent time segments. The particular choice of N_{ch} and w_{ol} results in a trade-off between frequency and time resolution (as well as information duplication) and must be carefully chosen by the user bearing in mind the typical characteristics of the wave, and data in question. We note in passing that our value of w_{ol} is not specified precisely but is given as a value with some tolerance, e.g., $w_{ol} = 30 \pm 1\%$, and an optimization algorithm chooses the precise overlap (for our choice of parameters) so as to fit as many time segments into a day's worth of data, minimizing the number of unused samples at the end of the (day's) time series.

[9] In Figure 1a we show the dynamic spectrogram of the X-component (geographic north) in greyscale, as a function of time and frequency for the first 12 hours (local time) of 06/06/2003. The rectangle bounding the region t = 2 to t = 5.4 hours, and f = 0.5 to f = 2.5 Hz contains a typical Pc1 pulsation which is analyzed further below (Figures 1e–1h). The vertical line at t = 2.79 hours (time segment i = 140) indicates a time period which will be used as an example to illustrate our peak detection algorithm (Figures 1b, 1d, 1f, and 1h).

[10] Labeling the windowed time signals at a specific time segment i as $x_i(t)$, $y_i(t)$, and $z_i(t)$, and the corresponding Fourier transforms $X_i(f)$, $Y_i(f)$, and $Z_i(f)$, we can define the covariance matrix in the frequency domain as

$$J^i(f) = \begin{vmatrix} X_i(f)X_i^*(f) & X_i(f)Y_i^*(f) & X_i(f)Z_i^*(f) \\ Y_i(f)X_i^*(f) & Y_i(f)Y_i^*(f) & Y_i(f)Z_i^*(f) \\ Z_i(f)X_i^*(f) & Z_i(f)Y_i^*(f) & Z_i(f)Z_i^*(f) \end{vmatrix} \quad (1)$$

where the asterisk superscript denotes complex conjugate. Using the off-diagonal elements of the covariance matrix (and noting that $|J^i_{lm}| = |J^i_{ml}|$ for $l, m = x, y, z$), we define the signal:

$$C^i(f) = |J^i_{12}|^2 + |J^i_{13}|^2 + |J^i_{23}|^2 \quad (2)$$

which provides a distribution of the total cross-covariance (squared) between all the components, as a function of frequency. The signal $C^i(f)$ is advantageous in that it is more immune to random noise than the autocovariance (diagonal) elements, since three mutually incoherent spatial signals will, by definition, have mutual coherencies of 0, resulting in a diagonal covariance matrix [*Means*, 1972]. The signal $C^i(f)$ is computed for every value of i to form a typical spectrogram representing the cross covariance power (Figure 1c).

[11] In both Figures 1a and 1c, the Pc1 pulsation (marked by the rectangle) is clearly visible by inspection since it

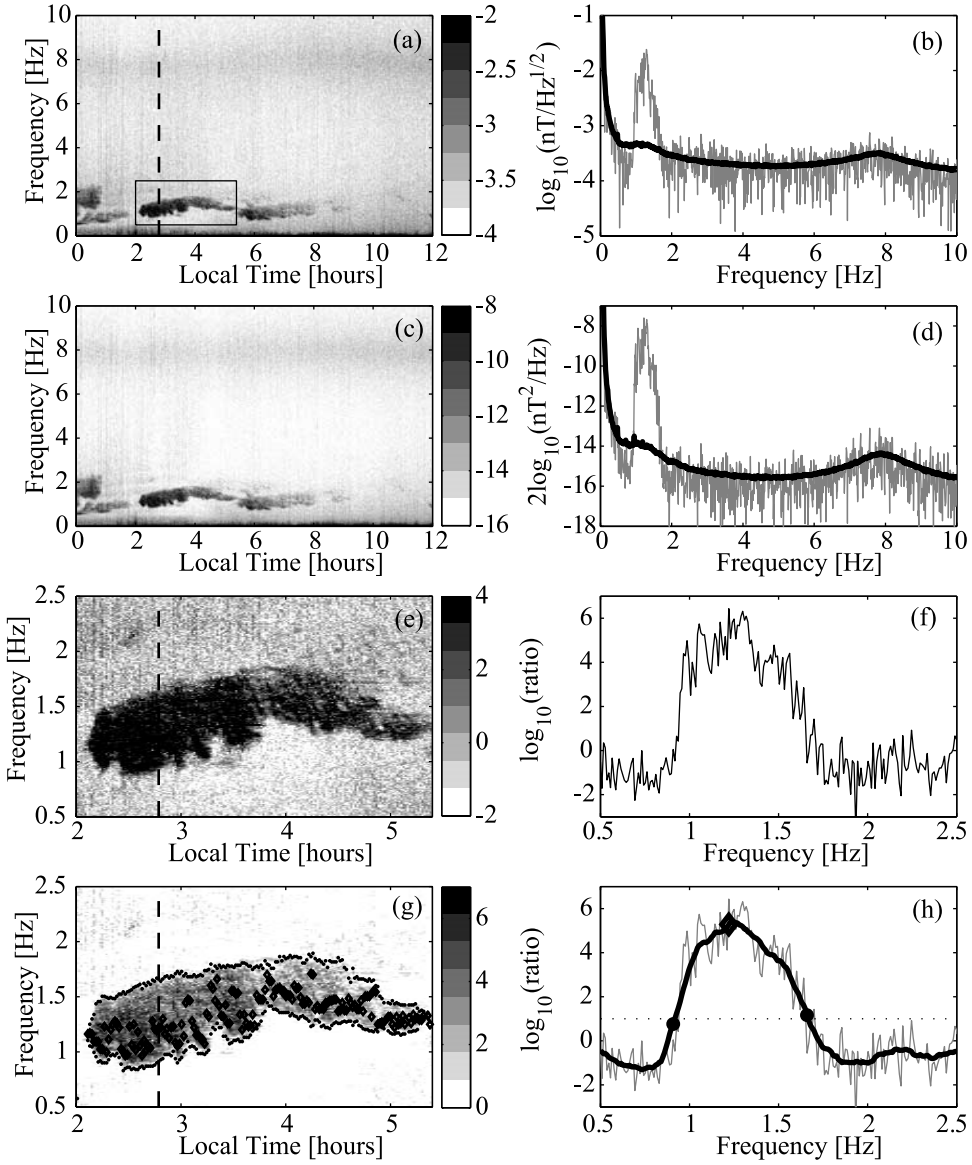

Figure 1. Frequency band identification. (a) Dynamic spectrogram of a single component of the field and (b) frequency spectrum at t = 2.79 hours (dashed line in Figure 1a) shown in gray, daily median shown as dark line. (c) Dynamic spectrogram of cross-covariance signal $C(f)$, (d) similar spectrum to Figure 1b, corresponding to Figure 1c. (e) Expanded portion of Figure 1c with daily median removed. Time limits indicated by rectangle in Figure 1a, (f) spectrum of Figure 1e at time of dashed line. (g) algorithm-identified Pc1 pulsation, and (h) as in Figure 1f showing sliding window averaged signal and threshold value at A_{th} = 1.

stands out sharply against the background. In order to automate the wave identification process, our algorithm needs to similarly estimate the background noise spectrum against which any unusual signals should be compared. This is achieved with a row-wise median extraction, i.e., if the spectrogram consists of 4096 rows (representing frequency) and 1205 columns (representing time segments), for every frequency component we select the median value of the 1205 elements so that we are left with 4096 values, representing the median value of the signal as a function of frequency, throughout that day. In Figures 1b and 1d the heavy line represents the daily median, $\log_{10}[M(f)]$, of Figures 1a and 1c, respectively, together with the signals at t = 2.79 hours shown as the thin gray lines ($\log_{10}[X_{140}(f)]$ and $\log_{10}[C^{140}(f)]$ in Figures 1a and 1c, respectively). The spectral peak of the Pc1 pulsation clearly rises above the median in both cases.

[12] In Figure 1e we show the expanded spectrogram corresponding to the rectangle in Figure 1a but with the median removed, such that at each time segment the signal $\log_{10}[C^i(f)] - \log_{10}[M(f)]$ is plotted. The dashed vertical line at t = 2.79 hours (as before) indicates the time segment at which the spectrum in the right panel (Figure 1f) is plotted. Note that the background signal is now distributed near unity (zero on the logarithmic scale in Figure 1f), while the spectral peak is several orders of magnitude more intense.

[13] As a final step, a sliding-average window is applied to the normalized signal and a threshold detection level is set. In our case, the sliding average was chosen to be $w_{\text{slide}} = 1\%$ of the sampling frequency (or 0.4 Hz) and the threshold value was set at $A_{th} = 1$ on the logarithmic scale, to ensure that detected signals were at least one order of magnitude greater than the background. The results of the sliding average and the original signal are shown in Figure 1h as the heavy line and light gray line, respectively, and the threshold level is shown as the dotted line. A spectral peak is detected when the averaged signal exceeds the threshold, resulting in three recorded frequencies: a bottom frequency ($f_{\text{bot}}^{\text{pk}} = 0.91$ Hz), a top frequency ($f_{\text{top}}^{\text{pk}} = 1.66$ Hz), and a frequency of maximum power ($f_{\text{max}}^{\text{pk}} = 1.22$ Hz), indicated in Figure 1h as two filled circles and a diamond symbol, respectively.

[14] It is also necessary to choose a lower and upper cutoff frequency, in our case $f_{\text{bot}}^{\text{cut}} = 0.1$ Hz and $f_{\text{top}}^{\text{cut}} = 10$ Hz, and a minimum width for the spectral peak, $\Delta f_{\text{peak}} = 0.1$ Hz. The spectral peak is considered valid when $f_{\text{bot}}^{\text{pk}} > f_{\text{bot}}^{\text{cut}}$, $f_{\text{top}}^{\text{pk}} < f_{\text{top}}^{\text{cut}}$, and $[f_{\text{top}}^{\text{pk}} - f^{\text{pk}}{}_{\text{bot}}] > \Delta f_{\text{peak}}$. In Figure 1g we show the same dynamic spectrogram as in Figure 1e, on a lighter color scale, and overlay the automatically identified series of $f_{\text{bot}}^{\text{pk}}$, $f_{\text{top}}^{\text{pk}}$, and $f_{\text{max}}^{\text{pk}}$ values. As shown, there is excellent agreement between visual and automatic identification.

[15] A few points should be noted at this stage: first, we choose to use a median value which is a function of frequency $M(f)$ because the background noise spectrum could (in some environments and/or frequency regimes) be changing very rapidly as a function of frequency, for example, $\sim\exp(f/f_0)$. Using only a single median value (not as a function of frequency) could cause the more intense parts of the noise spectrum (lower end in our example) to cross the threshold frequently, and the less intense parts of the spectrum (upper end in our example) to not cross the threshold, even when waves are present. For this reason we consider it vital to compare each frequency of the spectrum against the median value at that particular frequency. Second, if there is a long-enduring, constant frequency tone present in the data for most of the day, the daily median value will be set to the level of the tone, at the tone's frequency, and the tone itself will not be detected as a wave event. If a coincident Pc1 wave occurs over the frequency band covered by the tone, it will be detected only if it exceeds the tone's amplitude by a significant amount (set by A_{th}).

[16] In both cases discussed above, our detection algorithm was designed to closely mimic the way a human would manually detect wave activity by looking at a spectrogram, i.e., by looking for intense patches, where "intense" means that the power in the portion of the spectrum in question appears to be significantly larger than surrounding values, which could vary with frequency in some regular way. Constant tones that run across the length of the day would be rejected as legitimate wave events both by the human user, and our detection algorithm. However, in cases where the data block is on the order of the wave duration, the "data block median" may be set too high and the detection algorithm (and a human user with no prior experience) will not be able to detect the wave event. The solution in this case is to consider data blocks that are significantly longer than the typical duration of the wave events being sought in order to extract a meaningful median or to build some experience into our algorithm (to use the human user analogy) and calculate a median value over several data blocks. In the present case this problem is considered very unlikely since Pc1 pulsations exhibit a strong diurnal effect (which restrict them to durations of a few hours), and the frequency band exhibits some drift.

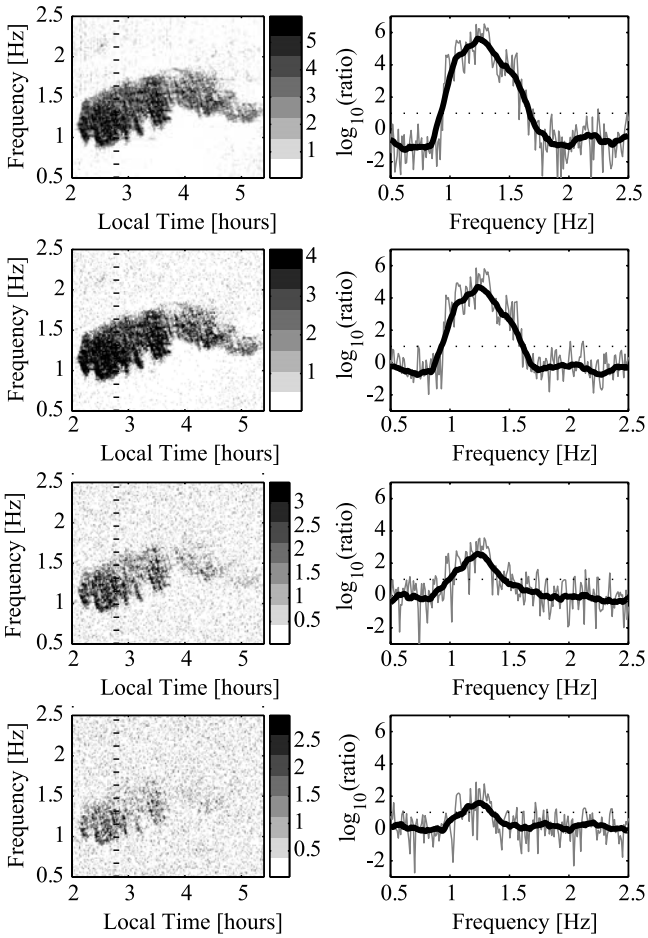

Figure 2. Spectral peak identification with noise. Columns correspond to Figures 1f and 1g, with each row having increasing noise, from top to bottom: SNR = 0, 0.1, 0.5, and 1.

[17] We examine the efficacy of our spectral peak detection technique in the presence of noise in Figure 2. The noise is injected by generating a normally distributed time series with a standard deviation $\sigma_{x,y,z}^{\text{noise}}$ equal to some fraction of the standard deviation of the corresponding time series $\sigma_{x,y,z}^{\text{sig}}$. The noise is added directly into the output signal of the analog-digital converter (i.e., before transfer function compensation) to simulate electronic noise coupling into the circuit. The rows of Figure 2 correspond to noise levels of $\sigma^{\text{noise}}/\sigma^{\text{sig}} = 0$, 0.1, 0.5, and 1. Columns 1 and 2 are computed in the same way as Figures 1f and 1h, except for the addition of the noise. Figure 2 illustrates that even though the spectral peak of the signal blends progressively into the noise floor as the level of noise is increased, for a low signal to noise ratio of 1 (bottom row) the signal is nevertheless detectable using our method.

[18] The robustness of our detection algorithm to noise is anticipated on theoretical grounds, since it can be shown

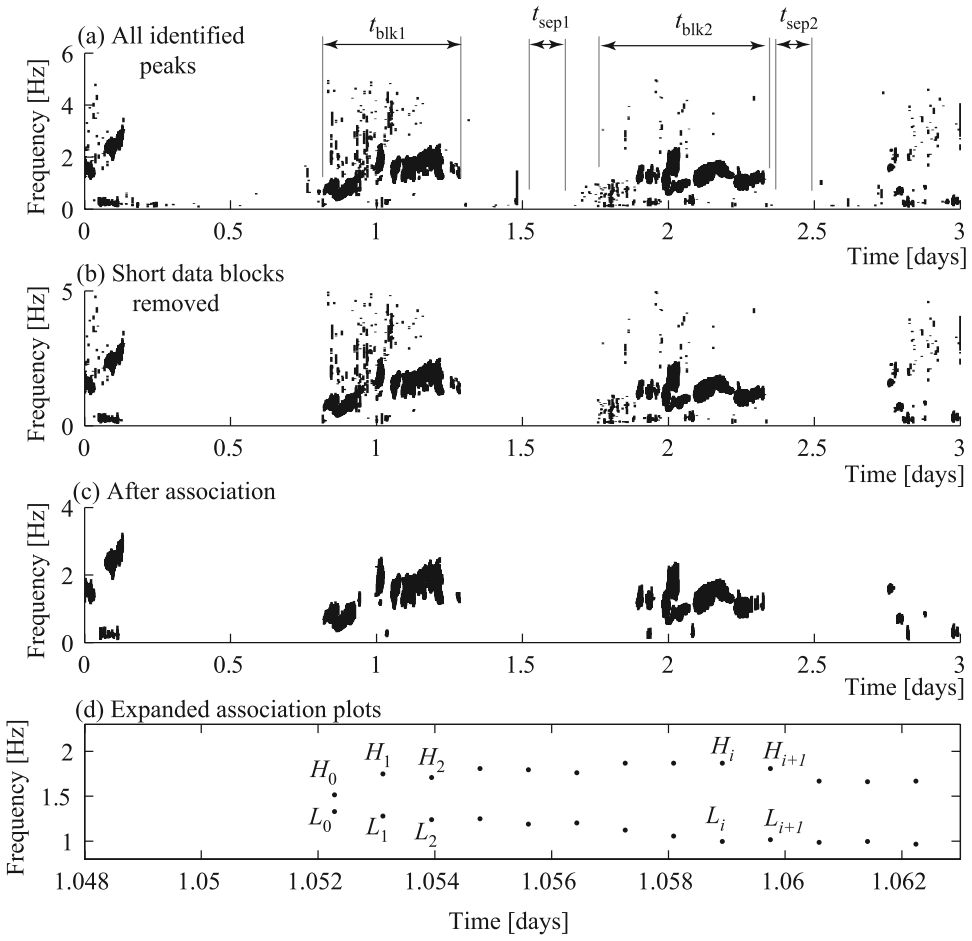

Figure 3. Illustration of spectral peak association; (a) all identified spectral peaks, (b) short data blocks removed, when $t_{blkn} < t_{blk}^{min}$, (c) all nonassociated spectral peaks removed, and (d) expanded portion of Figure 3c illustrating spectral peak association algorithm.

[*Means*, 1972] that incoherent noise is expected to enter only the diagonal elements of the covariance matrix and not the off-diagonal elements, which we use in our technique. However, in practice, when comparing detection efficiency using the diagonal elements of the covariance matrix against the detection efficiency using the off-diagonal elements of the covariance matrix (not shown), we have found only marginal improvement. Nonetheless, we use the off-diagonal elements in our identification algorithm since these terms provide direct information about wave properties (section 2.3) which we can use as a check on our temporal grouping technique (section 2.2).

[19] We have thus demonstrated that the spectral peak identification algorithm presented in this section works in the presence of severe noise and produces a list of three frequencies f_{bot}^{pk}, f_{top}^{pk}, and f_{max}^{pk} per identified peak as discussed above. Of course, any number of spectral peaks can occur in a single time segment, and the next task of the algorithm is to associate spectral peaks in contiguous time segments into individual, discrete wave events.

2.2. Temporal Grouping

[20] The temporal grouping of individual spectral peaks into discrete wave events proceeds as follows: (1) given a list of identified peaks, we search for extended quiet periods which are then used to divide the list into smaller data blocks which can be processed individually; (2) the data blocks are checked for length and those blocks that are shorter than a specified length are discarded; (3) each data block is then individually traversed beginning with the first spectral peak. Adjacent peaks are tested for continuity, and if a continuity occurs, the following peak is tested, and so on. When no continuity of spectral peaks is found over several adjacent time segments, the end of the wave event is signaled. In this way continuous wave events are identified and progressively removed from the data block and stored separately. Peaks that do not exhibit continuity are discarded until the entire data block is emptied, being divided essentially into either positively identified wave events or spurious peaks which are treated as noise.

[21] To illustrate our procedure, we show the list of identified spectral peaks in Figure 3a, where day 0 corresponds to local midnight on 4 June 2003 and day 3 is local midnight on 7 June 2003. The first half of day 2 (6 June 2003, 0000–1200) is the period previously shown in Figures 1 and 2. In Figure 3a the spectral peaks are drawn as vertical lines at every time segment, extending from f_{bot}^{pk} to f_{top}^{pk}, with no filtration, i.e., this is the raw output of the spectral identification procedure described in section 2.1. The list of peaks is then divided into contiguous data blocks

as shown by the labels t_{blk1}, t_{blk2}, etc., by looking for temporal separations (where no spectral peaks have been identified) in the data greater than a critical value, i.e., $t_{sepn} > t_{sep}^{min}$, where t_{sepn} is the nth temporal separation of the peaks in our data (e.g., t_{sep1} and t_{sep2} in Figure 3a), and t_{sep}^{min} = 10 min in our example. Since the grouping algorithm is memory-intensive, it is computationally much more efficient to divide the entire list of identified spectral peaks into smaller data blocks and perform spectral peak grouping on each data block in turn, than to perform it on the entire list. Using predefined time periods for processing (e.g., 1 day at a time or 1 week at a time) could achieve the same goal, but since the waves of interest could (and typically do) cross from one day to the next, or one time period to the next, they would be artificially split into two separate events (across the boundary) as opposed to one. We have thus chosen to split the data in a more natural way, by using extended quiet periods in our data to serve as the boundaries that separate individual data blocks (Figure 3a).

[22] The temporal separation described above results in a number of data blocks containing groups of spectral peaks. The next step is to discard data blocks that are too short, i.e., impulsive bursts. Figure 3b shows the results of this filtration, where we only retain data blocks with $t_{blkn} > t_{blk}^{min}$ and t_{blk}^{min} = 10 min in our example. As shown, many of the spurious spectral peaks have been eliminated, but those peaks that occur in conjunction with other large, contiguous blocks of data still remain.

[23] As a final step, we traverse each data block separately and check for association between adjacent spectral peaks, discarding those peaks that do not show association. The results of the association check are shown in Figure 3c, where we have now eliminated all spurious spectral peaks, and are left only with genuine Pc1 pulsation events (this is easily verified by inspection).

[24] The association procedure is illustrated in Figure 3d, where we show an expanded portion of Figure 3c between $1.048 < t < 1.063$ days and $0.8 < f < 2.3$ Hz, and the lines which marked spectral peaks in Figures 3a, 3b, and 3c are now replaced with points at f_{bot}^{pk} and f_{top}^{pk} at every time segment, as we did previously in Figure 1g. We use the shorthand symbols L and H to represent the low- and high-frequency bounds of the spectral peaks, respectively, and the subscript represents the spectral peak number in the data block. Beginning with the first peak (L_0, H_0), we test the next peak for continuity, such that the Boolean condition:

$$(L_{i+1} < H_i) \text{ AND } (H_{i+1} > L_i) = \text{TRUE} \tag{3}$$

is satisfied, ensuring nonzero spectral overlap between adjacent peaks. If (3) is satisfied, the following peak is tested, and so on until (3) is no longer satisfied. If a certain minimum number of peaks are found to be associated (i.e., (3) holds between any adjacent peaks), corresponding to a minimum duration t_{Pc1}^{min} then those peaks are grouped together into a "wave event," stored in a separate file and removed from the data block. If no association is found for a certain temporal peak, or association of several peaks is found, but the total duration of the event is less than t_{Pc1}^{min}, those peaks are assumed to be spurious noise and are discarded. In our case, we have chosen t_{Pc1}^{min} = 5 min. There is also the possibility that a spectral peak at a particular time segment has not been identified by the frequency band identification procedure (section 2.1), due to added noise or reduced signal amplitude, but the spectral peaks surrounding this time segment are indeed associated. To account for such cases, we extend the definition of "associated peaks" (or spectral overlap) from satisfying (3) at adjacent time segments, to that of satisfying:

$$(L_{i+j} < H_i) \text{ AND } (H_{i+j} > L_i) = \text{TRUE} \tag{4}$$

where j is incremented from 1 to some value such that $t_{i+j} - t_i \leq t_{Pc1}^{skip}$, and in our case we choose t_{Pc1}^{skip} = 3 min which allows a maximum of two time segments to be skipped.

[25] Each of the identified spectral peaks at time segment i is compared with the spectral peaks at time segment $i + 1$ (or $i + j$ in the more general case), starting from the lowest-frequency spectral peak and ending at the highest-frequency spectral peak. If an association is found (i.e., (4) is satisfied), the next time segment is checked for association by traversing all the identified spectral peaks from low- to high-frequency spectral peaks. The end of the event is marked when no associated spectral peaks have been detected in t_{Pc1}^{skip}, and all the entire event is stored in a separate file and removed from the data block. The process then continues until all the spectral peaks have been removed from the data block, as either spurious noise or wave events. In the case when two or more frequency bands merge into a single frequency band, the spectral overlap criterion alone will associate the lower-frequency band with the "merged" band, and remove this event from the data block, and the upper-frequency band will remain as an isolated event. However, one might wish to associate the frequency band, or individual spectral peaks, which are most "consistent" or "continuous" with adjacent peaks or bands when there is a choice. So far, our technique has been somewhat comparable to a simple edge detection algorithm, but due to the hierarchial way in which we have set up our problem and reduced our data, we are now able to compute the complete wave parameters for each wave event, which can be used in further filtering, continuity checking, or characterization, as discussed in the following section.

2.3. Continuity in Polarization Parameters

[26] The polarization parameters of a plane wave provide important information on the type of wave being studied, its origin, and the propagation characteristics of the medium through which it traveled. There have been a number of techniques published in the literature to determine the polarization parameters [e.g., *Born and Wolf*, 1970; *McPherron et al.*, 1972; *Means*, 1972; *Samson*, 1973; *Samson and Olsen*, 1980; *Rezeau et al.*, 1993; *Santolik et al.*, 2003]. The majority of these techniques use as their starting point the magnetic divergence equation $\mathbf{B}\cdot\mathbf{k} = \mathbf{0}$, which is multiplied by the conjugate of the magnetic signal $\mathbf{B}^*$, to give three mutually dependent equations:

$$\sum_{l=1}^{3} B_l B_m^* k_l = 0, \qquad l, m = x, y, z, \tag{5}$$

where k_l are the components of the wave normal vector, $B_l B_m^*$ form the elements of the Hermitian covariance matrix

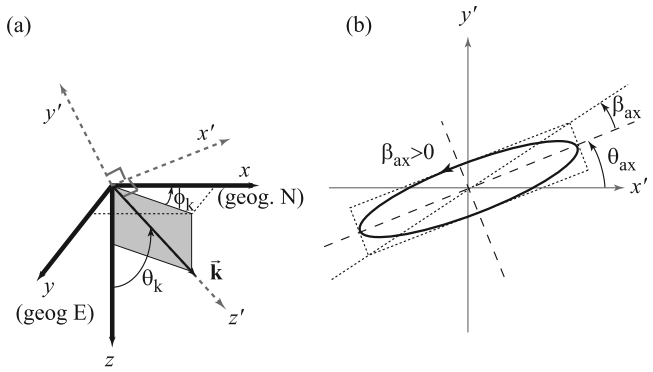

Figure 4. Orientation of axes. (a) rotation of XYZ coordinates into the principal coordinate system, (b) polarization ellipse in principal coordinates.

after appropriate manipulation, and we have now switched our notation from a general three-component vector, to a specific coordinate system to present our example Pc1 event, where the x, y, and z directions correspond to geographic north, geographic east, and vertically downward, respectively. In our case, we follow the methodology of *Means* [1972] due to its relative simplicity and use of the imaginary part of the off-diagonal elements of the covariance matrix, which is more immune to random noise than methods utilizing the real diagonal part of the covariance matrix [e.g., *McPherron et al.*, 1972].

[27] At a particular time segment i, we use the covariance matrix $J^i(f)$ (from (1)) and the bounding frequencies of a particular spectral peak $f_{\text{bot}}^{\text{pk}}$, $f_{\text{top}}^{\text{pk}}$, to obtain the band-integrated covariance matrix S:

$$S_{lm} = \int_{f_{\text{bot}}^{\text{pk}}}^{f_{\text{top}}^{\text{pk}}} J_{lm}^{i}(f)df \tag{6}$$

where l, $m = x, y, z$, and it is now understood that S is computed at a specific time segment i and over the identified frequency band.

[28] The wave normal vector is obtained directly from the imaginary part of S (S^I) [*Means*, 1972] as

$$k_x = \frac{S_{yz}^I}{a}, k_y = \frac{-S_{xz}^I}{a}, k_z = \frac{S_{xy}^I}{a}, \tag{7}$$

where $a^2 = S_{xy}^2 + S_{xz}^2 + S_{yz}^2$, and $k^2 = k_x^2 + k_y^2 + k_z^2 = 1$. The wave normal vector direction can be described by a set of polar angles (θ_k, ϕ_k), as

$$\theta_k = \arctan\left(\frac{\sqrt{k_x^2+k_y^2}}{k_z}\right),$$
$$\phi_k = \begin{cases} \arctan\left(k_y/k_x\right) & \text{for} \quad k_x \geq 0, \\ \arctan\left(k_y/k_x\right) - \pi & \text{for} \quad k_x < 0, k_y < 0, \\ \arctan\left(k_y/k_x\right) + \pi & \text{for} \quad k_x < 0, k_y \geq 0, \end{cases} \tag{8}$$

where we have been careful to locate ϕ_k in the correct quadrant [*Santolik et al.*, 2003]. Figure 4a illustrates the orientation of the **k**-vector and polar angles (θ_k, ϕ_k).

[29] The remainder of the procedure takes place in the principal coordinates, which requires that we rotate our coordinate system in such a way that the new z-axis is aligned with the **k**-vector. We use the general rotation matrix $R = B_R\, C_R D_R$, composed of three successive counterclockwise rotations of the axes about the z-, x-, and z-axes, respectively, defined by the Eulerian angles ϕ_R, θ_R, and ψ_R [*Goldstein*, 1965, p. 109], such that

$$B_R = \begin{bmatrix} \cos(\psi_R) & \sin(\psi_R) & 0 \\ -\sin(\psi_R) & \cos(\psi_R) & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{9}$$

$$C_R = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\theta_R) & \sin(\theta_R) \\ 0 & -\sin(\theta_R) & \cos(\theta_R) \end{bmatrix} \tag{10}$$

$$D_R = \begin{bmatrix} \cos(\phi_R) & \sin(\phi_R) & 0 \\ -\sin(\phi_R) & \cos(\phi_R) & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{11}$$

where $\phi_R = \phi_k - \pi/2$, $\theta_R = -\theta_k$, and $\psi_R = 0$, ensuring that the new z-axis (z') is parallel to the **k**-vector, x' is parallel to the horizontal plane (the former $x - y$ plane), and y' completes the RH coordinate set. The rotated coordinate system is illustrated in Figure 4a.

[30] Applying the similarity transformation to our band-integrated covariance matrix S, we obtain $S' = RSR^{-1}$, and note that only the upper [2×2] submatrix contains nonzero elements and is retained in further calculations. The remainder of the polarization parameters can now be obtained directly [*Fowler et al.*, 1967].

[31] The matrix S' is divided into a polarized and unpolarized part, P and U, respectively, given by

$$U = \begin{bmatrix} D & 0 \\ 0 & D \end{bmatrix} \tag{12}$$

where

$$D = \frac{1}{2}\left(S'_{xx} + S'_{yy}\right) - \frac{1}{2}\sqrt{\left(S'_{xx} + S'_{yy}\right)^2 - 4|S'|} \tag{13}$$

where $|S'|$ is the determinant of S', and $P = S' - U$.

[32] The polarization ratio, defined as the ratio of polarized power to total power is given by

$$R_{\text{pol}} = \sqrt{1 - \frac{4|S'|}{\left(S'_{xx} + S'_{yy}\right)^2}} \tag{14}$$

The angle between the major axis of the polarization ellipse and the x'-axis is defined by the angle θ_{ax} determined from

$$\tan(2\theta_{\text{ax}}) = \frac{2\Re e\{P_{xy}\}}{P_{xx} - P_{yy}} = A_{\text{ax}} \tag{15}$$

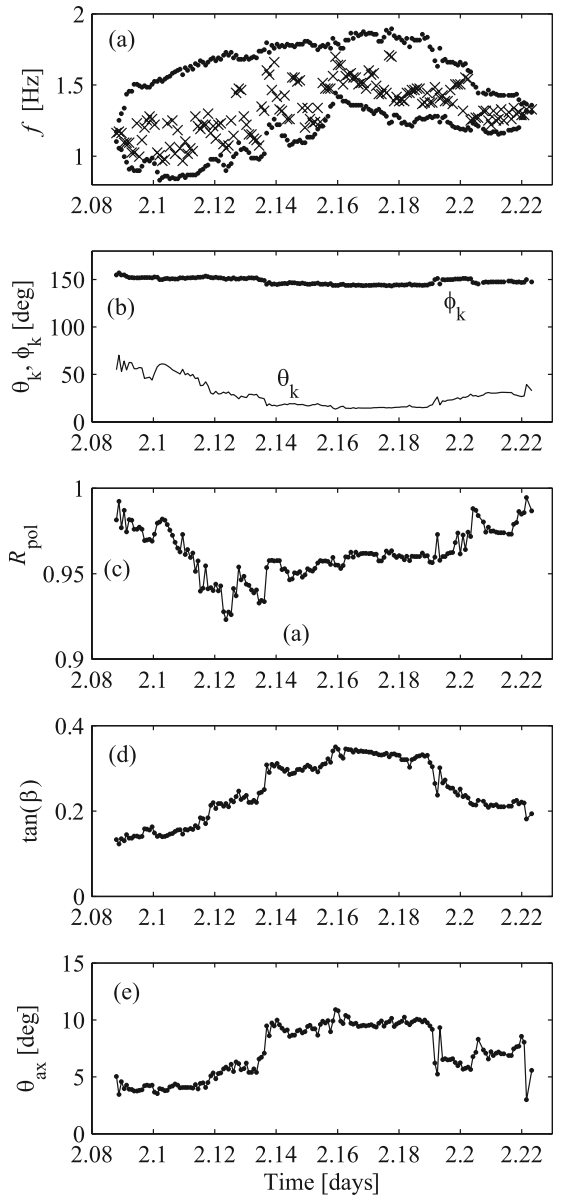

Figure 5. Polarization properties. (a) algorithm-identified wave event, (b) wave normal vector polar angles, (c) polarization ratio, (d) ellipticity, and (e) axis inclination angle.

giving

$$2\theta_{\rm ax} = \begin{cases} \arctan(A_{\rm ax}) & \text{for} \quad S'_{xx} > S'_{yy}, \\ \arctan(A_{\rm ax}) + \pi & \text{for} \quad S'_{xx} \le S'_{yy} \;\&\; A_{\rm ax} > 0, \\ \arctan(A_{\rm ax}) - \pi & \text{for} \quad S'_{xx} \le S'_{yy} \;\&\; A_{\rm ax} \le 0 \end{cases} \tag{16}$$

which gives the correct angle of rotation. To ensure that $\theta_{\rm ax}$ is in the range $(-\pi/2; \pi/2)$, we further add π if $\theta_{\rm ax} < -\pi/2$ and subtract π if $\theta_{\rm ax} > \pi/2$. The explicit adjustment of $\theta_{\rm ax}$ in (16) is necessitated by the fact that direct inversion of (15) gives a range of $-\pi/4 \le \theta_{\rm ax} \le \pi/4$, corresponding to the angle between the x'-axis and either the minor or the major axes of the polarization ellipse, which is insufficient for our purposes. The steps outlined above ensure that $\theta_{\rm ax}$ is indeed the angle between the x'-axis and the major axis of the polarization ellipse. As a final note, when $\theta_{\rm k} = 0$, the azimuthal angle $\phi_{\rm k}$ becomes arbitrary, so we set it to $\pi/2$, thus ensuring that $\theta_{\rm ax}$ returns an angle relative to the positive x-axis (geographic north).

[33] The ellipticity and sense of polarization are described by the angle $\beta_{\rm ax}$, where

$$\sin(2\beta_{\rm ax}) = \frac{i\left(P_{yx} - P_{xy}\right)}{\sqrt{\left(P_{xx} - P_{yy}\right)^2 + 4P_{yx}P_{xy}}} \tag{17}$$

and $\tan(\beta_{\rm ax})$ gives the ratio of the minor axis to the major axis (ellipticity), the sign of $\beta_{\rm ax}$ indicating the sense of polarization, $\beta_{\rm ax} > 0$ ($\beta_{\rm ax} < 0$) corresponding to RH (LH) rotation about the z'-axis. The coordinate transformation $(x, y, z) \rightarrow (x', y', z')$, and angles $\theta_{\rm ax}$ and $\beta_{\rm ax}$ are shown in Figures 4a and 4b, where the sense of rotation has been indicated for $\beta_{\rm ax} > 0$.

[34] Since the equation $\mathbf{B} \cdot \mathbf{k} = 0$ is being solved, the magnitude $k = |\mathbf{k}|$ can be multiplied by any arbitrary scalar constant (positive or negative) and hence does not contain any useful information, so it is normalized to unity for convenience. Consequently, there is an inherent ambiguity in the wave normal direction, since both $\mathbf{k}$ and $-\mathbf{k}$ are solutions. This ambiguity further reflects in the sign of $\beta_{\rm ax}$ as follows: suppose a circularly polarized wave is detected, with the phase front lying in the x-y plane, such that $\mathbf{k}$ lies along the z-axis and that when viewed from above, the sense of polarization is clockwise. If $\mathbf{k}$ is chosen to lie in the direction of $+z$ (down), the sense of polarization is RH and $\beta_{\rm ax} > 0$, whereas if $\mathbf{k}$ is chosen to lie in the direction of $-z$, the sense of polarization is LH and $\beta_{\rm ax} < 0$. Since the only information we have is the rotation of the wave's magnetic vector, and the plane in which it lies, both answers are legitimate solutions. This ambiguity can be resolved by either introducing additional information, such as an electric field vector which would determine the direction of $\mathbf{k}$ (by calculating the direction of the Poynting flux), or making an assumption about either the direction of $\mathbf{k}$ or sense of polarization (for example knowing that whistler-mode waves in space always rotate in a RH sense). In our case, we restrict $\mathbf{k}$ to be in the positive z half-space (i.e., $\mathbf{k}$ always faces toward the center of the Earth) under the assumption that the wave impinges onto the magnetometer from above, and thus the sense of polarization is allowed to be RH or LH, consistent with past work [e.g., *Summers and Fraser*, 1972]. Since our magnetometer is located in the northern hemisphere where the geomagnetic field is directed into the Earth, the sense of RH and LH polarizations are consistent with the definitions commonly used in plasma physics, which use the positive direction of the static magnetic field as the reference.

[35] In Figure 5 we show an example calculation of the polarization parameters for the Pc1 pulsation event shown previously in Figures 1–3. Figure 5a shows the identified spectral peaks as a set of upper- and lower-frequency values (dots) and a frequency of maximum power within each band ("x"-symbol), as a function of time in days (corresponding to the timescale used in Figure 3). Figure 5b shows the polar angles ($\theta_{\rm k}$, $\phi_{\rm k}$) describing the $\mathbf{k}$-vector orientation, indicating that the wave normal is within $\sim 30°$ of the vertical between $t = 2.12$ and $t = 2.22$ days, oriented at $\sim 150°$ to the

x-axis which corresponds roughly to a south-southeast direction.

[36] The polarization ratio in Figure 5c indicates that $R_{pol} > 90\%$ for the entire duration of the event, which is significantly higher than the polarization ratio measured away from the identified spectral peaks where $R_{pol} \sim 0$. Figure 5d shows that the ratio of minor axis to major axis was relatively small, with $0.1 < \tan(\beta_{ax}) < 0.4$, and remained relatively close to the x'-axis, with $5° < \theta_{ax} < 10°$ roughly, as shown in Figure 5e. The sign of β_{ax} indicates that the sense of rotation was RH, consistent with typical characteristics of Pc1 pulsations [e.g., *Jacobs*, 1970, pp. 19–32] which routinely exhibit both a RH and a LH sense of rotation.

[37] As described in the above section, a complete description of the wave's polarization parameters is readily obtained with only a few steps beyond the computation of the Fourier transformation, which is itself used in the production of standard spectrograms. The calculated wave parameters remain relatively constant for the duration of the event (see, e.g., ϕ_k) and can be used as an additional check on the temporal grouping algorithm (i.e., checking that the correct spectral peaks were grouped together). We may also choose to include a continuity criterion together with the spectral overlap criterion when we perform the temporal grouping, i.e., in addition to satisfying (4), we also require that $|(a_{i+j} - a_i)/a_i| < \chi$, where χ is some threshold value (e.g., 0.1 for less than 10% variation), and a_i represents any of the polarization parameters (e.g., θ_k, ϕ_k, etc.) at time segment i.

[38] We can also use any of the polarization parameters to perform additional filtering, for (1) quality, e.g., retaining only events with $R_{pol} > 0.8$, or (2) a particular wave characteristic, e.g., identifying only RH-polarized events, only events above a certain frequency threshold, only waves propagating in a given direction, and so on. Of course, we may choose to do no additional wave filtering beyond the temporal grouping algorithm (section 2.2), and simply study the morphology of all detected waves over some period of time. An example of the latter is presented below.

3. Case Study: June–December 2003

[39] To demonstrate our technique over an extended period, we use the Parkfield, California, search coil data (c.f. section 2) during the period 06/01/2003–12/31/2003, several examples of which are shown in Figures 1, 3, and 5. Our aim in this section is to present the typical outputs of our algorithm and possible uses of this technique on a typical data set, as opposed to conducting a study of the characteristics of the Pc1 waves themselves, which is too domain-specific and will not be handled in the current paper but deferred to future work.

[40] During the period 06/01/2003–12/31/2003, 1336 wave events were identified in the data using the parameters listed in Table 1. For each identified wave event, an average quantity was obtained as

$$\overline{a_n} = \frac{1}{T}\int_{t_{n0}}^{t_{n0}+T} a_n(t)dt \qquad (18)$$

Table 1. A Summary of the Values Which Must Be Set by the User in the Wave Identification Algorithm, Together With the Values Used in This Example

Parameter	Description	Value Used
N_{ch}	Number of samples in a time segment	4096
w_{ol}	Percentage overlap between time segments	30% ±1%
w_{slide}	Sliding-average window width (% of f_{samp})	1%
A_{th}	Threshold value for peak identification	1
f_{bot}^{cut}	Lower cutoff of spectral peaks	0.1 Hz
f_{top}^{cut}	Upper cutoff of spectral peaks	10 Hz
Δf_{peak}	Minimum bandwidth of spectral peak	0.1Hz
t_{sep}^{min}	Minimum separation time between data blocks	10 min
t_{blk}^{min}	Minimum duration of data block	10 min
t_{Pc1}^{min}	Minimum duration of Pc1 event	5 min
t_{Pc1}^{skip}	Maximum time to skip peaks within an event	3 min

where $a_n(t)$ is any of the wave properties (e.g., universal time t_{UT}, local time t_{LT}, $\Delta f = f_{max}^{pk} - f_{min}^{pk}$, θ_k, ϕ_k, R_{pol}, f_{max}^{pk}, etc.) associated with the n^{th} identified wave event, having a start time t_{n0} and duration T.

[41] Figure 6 shows the identified wave characteristics in the above period. In Figure 6a we plot f_{max}^{pk} versus $\overline{t_{UT}}$ (shown as day of year in 2003), such that each identified event is represented by a single point on the plot. Days 237–255 represent a data dropout and contain no identified events. This type of display essentially represents a summary of the dynamic spectrogram shown in Figure 3 and can be useful in identifying or associating coarse temporal features with other events. For example, we note the increase in occurrence and $\overline{f_{pk}^{max}}$ of Pc1 pulsations on days 326–329 (exceeding 4 Hz on day 328), following one of the largest ($D_{st} \sim -472$ nT) geomagnetic storms in the past 50 years on day 324 (20 November) [*Ebihara et al.*, 2005]. This behavior is consistent with past work which shows an increase in Pc1 occurrence and midfrequency with delay of 2–8 days after storms [*Tepley*, 1965].

[42] To obtain a sense for the distribution of various parameters, we show in the second row of Figure 6 histograms of a number of parameters. Figures 6a–6c show that the identified Pc1 events are predominantly observed at low frequencies (peak ~0.25 Hz; mean ~0.96 Hz), on the nightside, and at relatively high polarization ratios (peak ~0.87; mean ~0.73), respectively. The distribution of Pc1 events is consistent with past observations showing that the majority of Pc1 events at middle to low latitudes are observed on the nightside [*Campbell and Stiltner*, 1965; *Jacobs*, 1970, p.28]. The generation of Pc1 events (electromagnetic ion cyclotron waves, or EMIC) is believed to be driven by a cyclotron resonant interaction with anisotropic ring current protons [*Cornwall*, 1965] and are observed predominantly on the dayside at high L shells [*Fraser*, 1968; *Anderson et al.*, 1992a, 1992b; *Fraser and Nguyen*, 2001]. The EMIC waves then propagate along field lines predominantly in the left-hand mode, and couple into the ionosphere at high altitudes, where they mode-convert and

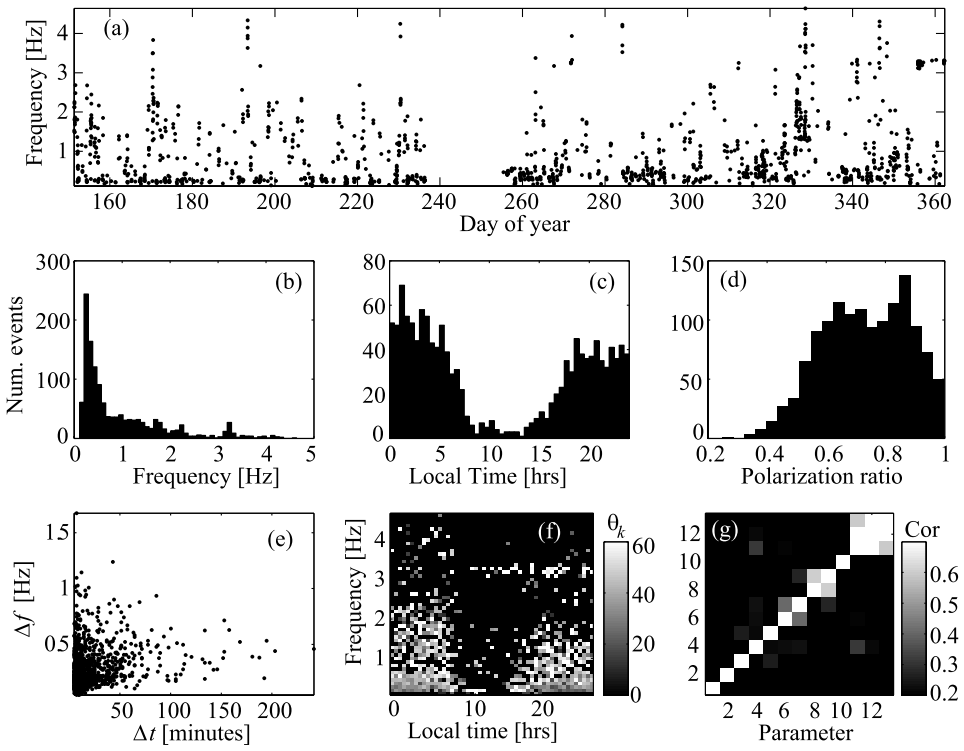

Figure 6. Statistical properties of identified wave events. (a) mean frequency versus mean time; distributions of (b) mean frequency, (c) local time distribution, (d) polarization ratio; (e) scatter plot of bandwidth versus duration; (f) wave normal zenith angle versus mean frequency and mean local time; (g) cross- correlation of 12 wave properties (see text for details).

travel predominantly in the right-hand mode in the ionospheric waveguide [*Fraser*, 1968; *Manchester*, 1968; *Jacobs*, 1970, p. 115]. The decreased occurrence of Pc1 events on the dayside at low latitudes is explained by the increased *E* layer absorption during the day, which strongly damps the EMIC waves as it propagates from its secondary source in the high-latitude ionosphere toward the equator [*Tepley*, 1965].

[43] The identified wave characteristics can be used to search for various associations as shown in the bottom row of Figure 6. Figure 6e is a scatterplot of the mean bandwidth ($\overline{\Delta f}$) of each identified Pc1 event, against event duration ($\overline{\Delta t}$), showing that the two parameters are weakly correlated. We compute the (nonparametric) Spearman rank order correlation coefficient [*Press et al.*, 2002, p. 640] to be $r_s \sim 0.313$, with a t-value of ~ 12.10 and P-value $< 10^{-16}$, indicating that the correlation is significant, and that the average event bandwidth tends to increase with event duration.

[44] Variation in three parameters is analyzed as shown in Figure 6f, where $\overline{\theta_k}$ is binned as a function of $\overline{f_{\max}^{pk}}$ and $\overline{t_{LT}}$. As shown in Figures 6b and 6c, the identified Pc1 events are clustered at low frequencies and around local midnight, but binning operation in Figure 6f reveals that there is also a systematic increase in $\overline{\theta_k}$ as a function of frequency at most local times, except for $t_{LT} = 1-4$, and $f > 3$ Hz, where $\overline{\theta_k}$ is again very low. The diurnal variation of $\overline{f_{\max}^{pk}}$ shows a maximum before dawn and a minimum just prior to dusk which is again consistent with past work [*Campbell and Stiltner*, 1965], and related to the variation in *F*-region characteristics from day to night.

[45] In Figure 6g we show a correlation matrix between 13 different quantities representing the mean characteristics of each identified wave event, respectively: (1) $\overline{t_{UT}}$, (2) $\overline{t_{LT}}$, (3) $\overline{\Delta t}$, (4) $\overline{f_{\max}^{pk}}$, (5) $\overline{\Delta f}$, (6) $\overline{R_{pol}}$, (7) $\overline{\theta_k}$, (8) $\overline{\phi k}$, (9) $\overline{\beta_{ax}}$, (10) peak magnitude, (11) K_p index, (12) plasmapause location (L_{pp}) estimated using the simple relation given by *Carpenter and Anderson* [1992], and (13) equatorial Helium gyrofrequency at L_{pp} (estimated using a centered dipole field model). This type of analysis is used to quickly identify correlations between a large number of parameters, which can then be used to infer various characteristics about the source of the waves, propagation characteristics, and so on. In our case the diagonals are perfectly correlated by definition, and the majority of the variables are uncorrelated. Moderate correlations exist between $\overline{\phi_k}$ and $\overline{\beta_{ax}}$ (variables 8 and 9), and $\overline{R_{pol}}$ and $\overline{\theta_k}$ (variables 6 and 7). Strong intercorrelations exist among variables 11–13 since these quantities are all calculated using K_p as input. As shown, there is not a strong correlation between the wave characteristics and instantaneous K_p value, which is consistent with past work [*Tepley*, 1965]. The wave frequency does not show strong correlation to the proton gyrofrequency at the plasmapause, which is likely related to the fact that while EMIC waves are clustered around the plasmapause, they can be generated well away from it [*Fraser and Nguyen*, 2001; *Meredith et al.*, 2003]. Analysis of further relations dealing with the detailed physics of Pc1 wave propagation is

considered beyond the scope of the present work and will be addressed in future studies.

4. Summary and Conclusions

[46] This paper discussed a new technique designed to automatically identify and characterize waves in three-axis data. This technique was demonstrated on a single Pc1 event recorded on a triaxial search coil magnetometer on 06/06/2003 and then applied to the 6-month period 06/01/2003–12/31/2003. The technique consists of three steps:

[47] 1. The first step is frequency band identification, in which we divide the time series data into short time segments, obtain the FFTs for each segment, and create a cross covariance signal $C^i(f)$ for each time segment using (2). The daily median value of $C^i(f)$ is then subtracted, the resulting curve smoothed, and all spectral peaks that exceed a given threshold are recorded.

[48] 2. The second step is temporal grouping, in which the identified spectral peaks are grouped into contiguous blocks and blocks shorter than a given threshold are removed. The remaining blocks are processed for temporal association, i.e., spectral peaks are grouped together if there is spectral overlap, equivalent to satisfying (4). If the number of grouped spectral peaks exceeds a minimum value, the peaks are recorded as an event.

[49] 3. The third step is polarization parameter calculation. Once individual events have been identified, the spectral peaks are used to calculate band-integrated polarization parameters such wave normal angles (θ_k, ϕ_k), polarization ratio R_{pol}, major axis angle θ_{ax}, and ellipticity and sense of polarization β_{ax}. These parameters can be used in further filtering, to check for continuity in a given event, or simply for characterization.

[50] The free parameters and their descriptions, as well as the chosen values used in this paper, are given in Table 1. In examining the 6-month period 06/01/2003–12/31/2003, 1336 events were identified which were clustered predominantly around low frequencies, on the nightside. Further analytical techniques were demonstrated for our example 6-month data set, which could be generally applied to a variety of data sets being analyzed with our technique.

[51] It should be noted that the parameters used in this paper were generic and only meant for illustrative purposes. In the current data set analyzed, our chosen parameters will favor the identification of unstructured pulsations but can be readily adjusted to identify a variety of other pulsations. For instance, to identify structured pulsations (hydromagnetic whistlers or "pearls"), the time segments should be chosen to be shorter and with a higher degree of overlap, to identify each "pearl" in the series individually. In fact, it is often the case that apparently structureless Pc1 pulsations do indeed have fine structure which is obscured by strong overloading signals or resolution limitations. The individual pearls will then need to be grouped with an additional layer of filtering, for example by looking for periodicity in the "center of mass" in the *f-t* plane.

[52] Similarly, our technique can be applied to a variety of frequency regimes and wave modes, measured both on the ground and on spacecraft, for example identification of lightning-generated whistlers, or magnetosheath lion roars which all exhibit spectral overlap between successive time segments for a given event. This technique is especially useful when the satellite data are rotated into some standard reference frame (for example, a field-aligned coordinate system) to facilitate the interpretation of the wave normal directions and other polarization parameters. In some instances, it is the "quiet band" between spectral peaks which is of interest, for example the helium stop band in electromagnetic ion cyclotron (EMIC) waves [*Mauk et al.*, 1981], which can be returned by our algorithm if two or more frequency bands are simultaneously present.

[53] If only one or two axis measurements are present, our technique can still be applied as follows: for the single axis measurements, only steps one and two will be performed, and the frequency bands will be identified using only signal power (as opposed to cross-covariance). If two axis measurements are present, the technique will proceed as in three dimensions, but all wave properties will be calculated in two dimensions assuming that the measurements are already in the principal coordinate system.

[54] As a final note, we mention that our technique is fairly general and can be tailored to analyze waves in a variety of situations. It has the distinct advantage (compared to simply looking for amplitude increases in a time series) that the waves are detected as a function of frequency and that the signal in which spectral peaks are sought is normalized by the median background frequency spectrum (c.f. section 2.1). This normalization implies, for example, that weak waves (e.g., at higher frequencies), which are nevertheless much stronger than the background in their frequency band, will be easily detected and not obscured by more intense signals (e.g., at lower frequencies) which are nevertheless weak compared to the background levels in their respective frequency band. The character of the background spectrum does not need to be known a priori because the algorithm computes it automatically when given sufficiently large data blocks (a day, in our case), which again underscores its generality.

[55] **Acknowledgments.** This work was supported by QuakeFinder LLC., Stellar Solutions Inc., the California Space Authority (CSA), and NASA grant NNG04GD16A. Magnetometer data was obtained from the Northern California Earthquake Data Center (NCEDC), contributed by the Berkeley Seismological Laboratory, University of California, Berkeley. The authors would like to thank Celeste V. Ford for her ongoing help and support. JB would like to thank Eftyhia Zesta, Peter Chi, and Richard M. Thorne for helpful discussions.

[56] Zuyin Pu thanks Timo Asikainen and Brian Fraser for their assistance in evaluating this paper.

References

Anderson, B. J., R. E. Erlandson, and L. J. Zanetti (1992a), A statistical study of Pc 1–2 magnetic pulsations in the equatorial magnetosphere: 2. Wave properties, *J. Geophys. Res.*, *97*, 3089.

Anderson, B. J., R. E. Erlandson, and L. J. Zanetti (1992b), A statistical study of Pc 1–2 magnetic pulsations in the equatorial magnetosphere: 1. Equatorial occurrence distribution, *J. Geophys. Res.*, *97*, 3075.

Born, M., and E. Wolf (1970), *Principles of Optics*, 4th ed., pp. 544–558, Elsevier, New York.

Bracewell, R. N. (2000), *The Fourier Transform and its Applications*, 3rd ed., McGraw-Hill, New York.

Campbell, W. H., and E. C. Stiltner (1965), Some characteristics of geomagnetic pulsations at frequencies near 1 c/s, *Radio Sci. J. Res.*, *69D*(8), 1117.

Canny, J. (1986), A computational approach to edge detection, *IEEE Trans. Pattern Anal. Mach. Intel.*, *8*, 679.

Carpenter, D. L., and R. R. Anderson (1992), An ISEE/Whistler model of equatorial electron density in the magnetosphere, *J. Geophys. Res.*, *97*, 1097.

Cornwall, J. M. (1965), Cyclotron instabilities and electromagnetic emission in the Ultra Low Frequency and Very Low Frequency ranges, *J. Geophys. Res.*, *70*, 61.

Ebihara, Y., M.-C. Fok, S. Sazykin, M. F. Thomsen, M. R. Hairston, D. S. Evans, F. J. Rich, and M. Ejiri (2005), ring current and the magnetosphere-ionosphere coupling during the superstorm of 20 November 2003, *J. Geophys. Res.*, *110*, A09S22, doi:10.1029/2004JA010924.

Erlandson, R. E., and B. J. Anderson (1996), Pc 1 waves in the ionosphere: A statistical study, *J. Geophys. Res.*, *101*, 7843.

Fowler, R. A., B. J. Kotick, and R. D. Elliott (1967), Polarization analysis of natural and artificially induced geomagnetic micropulsations, *J. Geophys. Res.*, *72*, 2871.

Fraser, B. J. (1968), Temporal variations in Pc1 geomagnetic micropulsations, *Plant. Space Sci.*, *16*, 111.

Fraser, B. J., and T. S. Nguyen (2001), Is the plasmapause a preferred source region of electromagnetic ion cyclotron waves in the magnetosphere?, *J. Atmos. Sol. Terr. Phys.*, *63*, 1225.

Goldstein, H. (1965), *Classical Mechanics*, Addison-Wesley, Boston, Mass.

Gurnett, D. A., and U. S. Inan (1988), Plasma wave observations with the Dynamics Explorer 1 spacecraft, *Rev. Geophys.*, *26*, 285.

Hamar, D., and G. Tarcsai (1982), High resolution frequency time analysis of whistlers using digital matched filtering, part 1: theory and simulation studies, *Ann. Geophys.*, *38*, 119.

Jacobs, J. A. (1970), *Geomagnetic Micropulsations*, Springer, New York.

Loto'aniu, T. M., B. J. Fraser, and C. L. Waters (2005), Propagation of electromagnetic ion cyclotron wave energy in the magnetosphere, *J. Geophys. Res.*, *110*, A07214, doi:10.1029/2004JA010816.

Manchester, R. N. (1968), Correction of Pc1 micropulsations at spaced stations, *J. Geophys. Res.*, *73*, 3549.

Mauk, B. H., C. E. McIlwain, and R. L. McPherron (1981), Helium cyclotron resonance within the Earth's magnetosphere, *Geophys. Res. Lett.*, *8*, 103.

McPherron, R. L., C. T. Russell, and P. J. Coleman Jr. (1972), Fluctuating magnetic fields in the magnetosphere, II. ULF waves, *Space Sci. Rev.*, *13*, 411.

Means, J. D. (1972), Use of the three-dimensional covariance matrix in analyzing the polarization properties of plane waves, *J. Geophys. Res.*, *77*, 5551.

Meredith, N. P., R. M. Thorne, R. B. Horne, D. Summers, B. J. Fraser, and R. R. Anderson (2003), Statistical analysis of relativistic electron energies for cyclotron resonance with EMIC waves observed on CRRES, *J. Geophys. Res.*, *108*(A6), 1250, doi:10.1029/2002JA009700.

Moore, G. E. (1965), Cramming more components onto integrates circuits, *Electronics*, *38*, 8.

Mursula, K., J. Kangas, and J. Kultima (1994), Looking back at the early years of Pc1 pulsation research, *Eos Trans. AGU*, *75*(31), 357–365.

Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery (2002), *Numerical Recipes in C*, Cambridge Univ. Press, New York.

Rezeau, L., A. Rouz, and C. T. Russell (1993), Characterization of small-scale structures at the magnetopause from ISEE measurements, *J. Geophys. Res.*, *98*, 179.

Samson, J. C. (1973), Descriptions of the polarization states of vector processes: Applications to ULF magnetic fields, *Geophys. J. R. Astron. Soc.*, *34*, 403.

Samson, J. C. (1991), Geomagnetic pulsations and plasma waves in the Earth's magnetosphere, in *Geomagnetism*, vol. 4, edited by J. A. Jacobs, p. 481, Elsevier, New York.

Samson, J. C., and J. V. Olsen (1980), Some comments on the decriptions of the polarization states of waves, *Geophys. J. R. Astron. Soc.*, *61*, 115.

Santolik, O., M. Parrot, and F. Lefeuvre (2003), Singular value decomposition methods for wave propagation analysis, *Radio Sci.*, *38*(1), 1010, doi:10.1029/2000RS002523.

Singh, R. P., D. K. Singh, A. K. Singh, D. Hamar, and J. Lichtenberger (1999), Application of matched filtering and parameter estimation technique to low latitude whistlers, *J. Atmos. Sol. Terr. Phys.*, *61*, 1081.

Stix, T. H. (1992), *Waves in Plasmas*, Springer, New York.

Summers, W. R., and B. J. Fraser (1972), Polarization properties of Pc1 micropulsations at low latitudes, *Planet. Space Sci.*, *20*, 1323.

Tepley, L. (1965), Regular oscillations near 1 c/s observed at middle and low latitudes, *Radio Sci. J. Res.*, *69D*(8), 1089.

Walker, R. J., T. A. King, and S. P. Joy (2005), Future directions in space physics data management, *Eos Trans. AGU*, *86*(18), Jt. Assem. Suppl., Abstract SH42A-01.

T. E. Bleier and C. Dunson, QuakeFinder LLC., 250 Cambridge Ave., Suite 204, Palo Alto, CA 94305, USA. (tbleier@quakefinder.com; cdunson@quakefinder.com)

J. Bortnik, Department of Atmospheric and Oceanic Sciences, University of California, Los Angeles, Room 7115 Math Sciences Bldg., Los Angeles, CA 90095-1565, USA. (jbortnik@gmail.com)

J. W. Cutler, Department of Aeronautics and Astronautics, Space Systems Development Laboratory, Stanford University, Room 269 Durand Bldg., 496 Lomita Mall, Stanford, CA 94305-4035, USA. (jwc@stanford.edu)